John Kinnear

Report Series
No 13

The Trends in Data Processing Costs

May 1980

The Butler Cox Foundation

The Butler Cox Foundation

The Butler Cox Foundation is a research group that examines major developments in the fields of computers, telecommunications and office automation on behalf of its subscribing members. The Foundation provides a set of 'eyes and ears' on the world for the systems departments of some of Europe's largest organisations.

The Foundation collects its information through its office in London and also through its associated offices in Europe and the US. It transmits its findings to its members in three main ways:

— Through regular written reports that give detailed findings and substantiating evidence.

— Through management conferences for management services directors and their senior colleagues, where the emphasis is on the policy implications of the subjects studied.

— Through professional and technical seminars where the members' own specialist managers and technicians meet with the Foundation research teams to review their findings in depth.

The Foundation is controlled by a Management Board whose members include representatives from the Foundation member organisations. The responsibilities of the Management Board include selecting topics for research and approving the Foundation's annual report and accounts, which show how the subscribed research funds have been employed.

The Management Board

David Butler	Chairman, Butler Cox & Partners Limited, Morley House, 26-30 Holborn Viaduct, London EC1A 2BP Tel: 01-583 9381
Michael Colin	Director of Advanced Systems, BL Systems Limited, CDC, PO Box 5, Redditch, Worcestershire B98 9HB Tel: (0527) 60133
George Cox	Managing Director, Butler Cox & Partners Limited (as above)
Martin Ray	Director of the Butler Cox Foundation, Butler Cox & Partners Limited (as above)
Bill Sherratt	Information and Computing Services Division, Shell UK Limited, PO Box 148, Shell-Mex House, Strand, London WC2R 0DX Tel: 01-438 3000
Bob Tester	BOC Limited, Hammersmith House, Hammersmith Bridge Road, London W6 9DX Tel: 01-748 2020

Publications 1977-80

Reports 1977-80

1/July 1977	Developments in data networks
2/August 1977	Display word processors
3/November 1977	Terminal compatibility
4/December 1977	Trends in office automation technologies
5/February 1978	The convergence of technologies
6/April 1978	Viewdata *(Out of print)*
7/May 1978	Public data services
8/June 1978	Project management
9/July 1978	The selection of a computerised PABX
10/November 1978	Public on-line information retrieval services
11/February 1979	Improving systems' productivity
12/June 1979	Trends in database management systems
13/April 1980	Cost trends in data processing
14/September 1979	The competitive equipment market
15/September 1979	Management services and the microprocessor
16/November 1979	Mainframe computers in the 1980s
17/February 1980	Electronic mail
18/April 1980	Distributed processing: management issues
19/June 1980	Office systems strategy
20/August 1980	The interface between people and equipment
21/October 1980	Private communications networks

Conference Transcripts 1977-80

1/May 1977 London	Computers, communications and office automation: the emerging pattern of information systems
2/November 1977 Stratford-upon-Avon	The management policy issues arising from emerging technologies
3/April 1978 London	Key issues in data processing
4/October 1978 Bristol	Management priorities for the 1980s
5/May 1979 Westchester, New York	Office automation: when and how will it happen?. Distributed processing: the policy issues for management services
6/November 1979 Birmingham	The changing role of management services
7/May 1980 Venice	System development — prospects for the 1980s
8/December 1980	New applications, new technologies

Abstract

Report Series No 13

The Trends in Data Processing Costs

by Edward Goldblum

May 1980

The dramatic and continuing developments in digital electronics have led to ever-cheaper computer components, but, despite this, most users continue to pay more every year for their data processing. This apparent paradox results from a combination of historical and technological factors that have shaped today's computer industry.

This report explores in detail the quantitative changes in data processing costs that have taken place over the last twenty years. These costs are divided, in the report, into five main categories: processors and memory, auxiliary storage devices, other peripheral equipment, services and facilities, and salaries and overheads. Each category is examined separately, and a great deal of supporting data is presented in the form of graphs and tables.

The report then analyses total system costs in terms of their overall size, growth rates, and composition. Several important conclusions regarding the interactions of users, manufacturers, and costs are drawn from this data.

The Butler Cox Foundation is a research group that examines major developments in the fields of computers, telecommunications and office automation on behalf of its subscribing members. The Foundation provides a set of "eyes and ears" on the world for the systems departments of some of Europe's largest organisations.

The Foundation collects its information through its office in London and also through its associated offices in Europe and the US. It transmits its findings to its members in three main ways:

— Through regular *written reports* that give detailed findings and substantiating evidence.

— Through *management conferences* for management services directors and their senior colleagues, where the emphasis is on the policy implications of the subjects studied.

— Through *professional and technical seminars* where the members' own specialist managers and technicians meet with the Foundation research teams to review their findings in depth.

The Foundation is controlled by a Management Board whose members include representatives from the Foundation member organisations. The responsibilities of the Management Board include selecting topics for research and approving the Foundation's annual report and accounts, which show how the subscribed research funds have been employed.

Report Series No. 13

THE TRENDS IN DATA PROCESSING COSTS

by Edward Goldblum

May 1980

Butler Cox & Partners Limited Morley House 26 Holborn Viaduct London EC1A 2BP

CONTENTS

1. INTRODUCTION 1

THE PURPOSE OF THIS REPORT 1
THE COMPONENTS OF COST 1
THE SCARCITY OF GOOD INFORMATION 2
GUIDELINES FOR DATA COLLECTION 3

2. PROCESSORS AND MEMORY 5

THE FACTORS THAT INFLUENCE PROCESSOR COSTS 5
TRENDS IN PROCESSOR SPEEDS AND PRICES 7
THE MARKET FOR MAINFRAME COMPUTERS 10
THE MARKET FOR MINICOMPUTERS 15
MEMORY COST TRENDS 15
THE MARKET FOR COMPUTER MEMORY 19
EXPENDITURE ON PROCESSORS AND MEMORY 19
SUMMARY 26

3. AUXILIARY STORAGE DEVICES 28

TRENDS IN DISC STORAGE 28
TRENDS IN MAGNETIC TAPE 31
THE UTILISATION OF DISCS AND TAPES 31

4. OTHER PERIPHERALS 36

LINE PRINTERS 36
UNIT-RECORD EQUIPMENT 39
DATA ENTRY EQUIPMENT 39
TERMINALS 43

5. SALARIES AND GENERAL OVERHEADS 45

UK AND US SALARY LEVELS 45
SALARY COSTS AND STAFFING LEVELS 46
GENERAL OVERHEADS 49

6. SERVICES AND FACILITIES 52

EQUIPMENT MAINTENANCE 52
PHYSICAL FACILITIES 54
DATA PROCESSING SUPPLIES 54
PURCHASED SOFTWARE AND SERVICES 56
DATA COMMUNICATIONS SERVICES 58

7. TOTAL SYSTEM COSTS 62

THE COMPONENTS OF TOTAL COST 62
GROWTH OF DP COSTS 62
THE CONCLUSIONS THAT WE DRAW 67

8. CONCLUSION 70

BIBLIOGRAPHY 71

CHAPTER 1

INTRODUCTION

THE PURPOSE OF THIS REPORT

In many organisations, management holds two views of expenditure on computer-based systems, which at first sight appear to be contradictory. These views are that:

1. Technological advance is constantly making computers cheaper.

2. Expenditure on computing is relentlessly increasing.

The computing industry (and more recently the telecommunications and the office equipment industries) are experiencing intense pressures in both technology development and product pricing. These pressures have already caused some established companies to move out of a particular sector of the industry and, in some extreme cases, out of business altogether. Others have been obliged to seek transfusions of external capital.

Nevertheless, it is rare to encounter an organisation that claims to have held constant (let alone to have reduced) its overall expenditure on systems. Even organisations that have contained or have reduced their data processing staff numbers do not make that claim.

A complete examination of the patterns of investment in information systems needs to cover a number of separate, but inter-related issues. The particularly important ones seem to be:

- The cost trends of equipment, services and staff.
- The patterns of expenditure on equipment, services and staff.
- The trends in benefits derived from the use of computer-based systems.
- Attitudes and approaches to making investment decisions, particularly in application areas where tangible cost displacement is not a sufficient justification.

This report concentrates on the first two of these issues, and it describes and analyses the trends in these areas. It concentrates on data processing, and it reports costs and expenditure in other areas (such as telecommunications) only where these costs impact upon data processing. The analysis looks back as far as twenty years to help identify genuine trends, as distinct from short-term fads and responses to transient market forces.

Everyone is aware that computing costs are changing, but this report aims to explain the way in which those costs are changing, by what proportions, and with what consequences for the computer user.

THE COMPONENTS OF COST

Total DP costs can be broken down for analysis in many different ways. Each such breakdown reflects a slightly different view of data processing (for example, functional, administrative, financial, etc.), and each is also fairly arbitrary. We have chosen to analyse costs as nearly as possible on a "like-with-like" basis, and this choice has led us to the following breakdown of system cost components, which we use in this report:

1. Processors and central memory
   These comprise arithmetic/logic units and their directly-addressable main memories in which programs and data normally reside during program execution. This category includes both mainframe computers and minicomputers.

2. Auxiliary storage devices
   These comprise magnetic data storage media of all kinds, such as tapes, discs, drums, and variants of these.

3. Other peripherals and hardware
   These comprise a wide variety of computer peripherals, notably printers, unit-record (card) equipment, terminals, visual display units, data entry equipment and communications devices.

4. Services and facilities
   These comprise accommodation, maintenance, data communications services, supplies, purchased software, and bureau services.

5. Personnel
   These comprise analysts and programmers, computer operators, and data entry staff.

In addition, the report includes a separate section on total system costs, their patterns of change, and their implications for data processing in large organisations.

THE SCARCITY OF GOOD INFORMATION

Anyone who sets out to survey historical trends in the computer industry immediately discovers that there is no single amalgamated, normalised body of accurate statistical data on the industry or its component parts. This deficiency is rather ironic, bearing in mind that the computer itself is widely used as a repository of precise and timely statistics on many subjects.

Reliable statistics on the computer industry are difficult to obtain for at least the following three reasons:

1. There is no generally-agreed set of definitions of even the most basic terms in computing. Thus, "minicomputers" may edge further and further into the "mainframe" or general-purpose category at one end of the spectrum, whereas, at the other extreme, a microcomputer (depending upon its use) may not be counted as a computer at all. Each manufacturer tends to use his own definitions and terms, and he does not make his choice with a view to making life easier for those who compile statistics.

2. Most manufacturers (and many users) are highly secretive about their activities and do not publish detailed statistics about their own operations. With some manufacturers, this secrecy stems from an understandable commercial reticence, but with others it is perhaps more motivated by a defensive wish to conceal their failures from public view. What this always means, however, is that any data required must be painstakingly extracted from advertisements, public announcements, private and confidential sources and rumours. Because all of these sources are notoriously unreliable, they can easily produce misleading conclusions. One reflection of this problem is that a thriving sub-industry has grown up around the single task of monitoring the activities of IBM.

3. Those organisations that do attempt to trace movements in the marketplace are severely handicapped by the industry's very rapid growth rate. Unlike many other industries, the growth of the computer industry cannot be accurately estimated merely by adding a fixed percentage increase to the previous year's results. The influx of new competitors into the market each year, and the departure of older ones, mean that traditional sources of data may quickly become inaccurate. New product lines are constantly being announced, and many of these lines have a very brief run in the market. All these factors introduce potential errors for those who seek to measure the underlying trends.

GUIDELINES FOR DATA COLLECTION

Because of these difficulties, we adopted the following approach in collecting the data which forms the basis of this report.

1. Using as few different sources as possible
   There are several commercial sources of regular and reasonably consistent data on the computer industry and marketplace, but we have used as few of these separate sources as possible. Independent sources rarely agree with one another to an accuracy of better than about 20%, and so it is not sensible to expect "raw" data to be more accurate than this. However, this report is more concerned with trends than with year-on-year absolute values, and consequently it is *consistency of data* that is needed most. By limiting the sources, we believe that the incidence of inconsistent errors has been minimised.

   We used the following principal sources of statistical data for this report:

   — International Data Corporation (IDC) of Waltham, Massachusetts, USA, whose *EDP Industry Report* is based on a very comprehensive set of well-maintained data on more than 45,000 computers in the USA.

   — Auerbach, Inc., of Philadelphia, Pa., USA, who publish extensive surveys of products in the DP industry.

   — Datapro Research Corporation, of Delran, New Jersey, USA, who publish reports similar to those of Auerbach.

   By far the most useful source of analysed data that we used was Montgomery Phister's monumental work, *Data Processing Technology and Economics* and its later supplement. This book draws mainly upon the same sources as those named above, but it includes original research material and analysis, as well as comments and interpretations of the raw data that show considerable insight. This work is recommended to anyone who is interested in gaining a comprehensive, quantitative view of data processing and its historical development.

2. Using American market data unless it is inappropriate
   The US computer market remains the largest in the world by far, accounting as it does for about 60% of the entire world computer market by value. The US computer industry supplies (either directly or indirectly) at least 90% of the world's computer systems. The US domestic market is also the most extensively monitored and thoroughly documented computer market.

   This is not to suggest that non-American computer users will have had the same experiences as their US counterparts. American patterns of computer usage are certainly different from those of (say) Italy or Australia. But non-US users probably differ from one another as much as they differ from Americans. No "non-US computer user profile" has yet emerged. Furthermore, many European users, as a result of having read so much literature about American users, have already established a view of where they stand in relation to users in the US.

American statistics have been used throughout this report partly for the reasons discussed above, partly because there is so much data available on the US market, and partly because it is not easy (if in fact it is possible) to find comparable data on other markets. Unfortunately, no statistics on installations that use European computers were available from the US sources. Because of this dominance of American statistics, and because of the absence of statistics relating to European computer manufacturers, some readers will undoubtedly wish to make allowances and compensations when relating the statistics in this report to their own installations.

3. Using US dollar prices and costs throughout
   The reasons that convinced us to use the American currency unit throughout a report written largely for European readers are the following:

   — As already mentioned, over 90% of all computer equipment is produced by US-based companies. About 55% of all computer hardware physically originates in the US. The dollar is therefore the "natural" accounting unit for comparing computer expenditure.

   — Historically, the US price index (currency deflator) has been more stable than those of many European countries. Therefore, the increases in money amounts due solely to inflation are less intrusive when the American currency is used.

   — Converting data from one currency to another realistically, over a period of twenty years, is extremely difficult. The time-varying inter-currency fluctuations since 1971 have been so large and erratic as to distort completely the value of items expressed in any currency other than the originating one.

   — Shipping costs, local taxation, and differential pricing policies of multinational companies all tend to distort the prices of items that originated at a single price in the US.

   With two exceptions, we have therefore adopted *historical US dollar amounts* (that is, the original amounts, *not* corrected for inflation) as the basic monetary measures for this report. The two exceptions are salaries and data communications costs, for which we use British costs expressed in sterling.

   It may seem unwise to ignore the effects of inflation in this way, given that inflation dominates financial analysis at present. It must be remembered, however, that inflation does not affect everyone equally, and there is no way of knowing exactly how a particular company accounts for the effects of inflation. The consumer price index, the growth of the money supply, the bank rate, the cost of corporate borrowing, and other indices all have their particular applications and their staunch advocates.

   In providing the original US dollar amounts, uncorrected for inflation, we expect that readers will make whatever year-by-year adjustments to these values they deem appropriate to their circumstances.

CHAPTER 2

PROCESSORS AND MEMORY

Irrespective of price movements, processors have always been the centre of computer systems and the focus of most evaluations and comparisons. This emphasis is not wholly appropriate, because most users today run applications programs that are primarily constrained by other factors than processor performance alone. Nevertheless, computer manufacuturers continue to base their systems around particular processor architectures, and some potential buyers still prepare to shop for new systems by arming themselves with the latest processor performance figures.

Such an approach may not be especially fruitful, but at least it has the advantage of being quantitative, and it allows a certain amount of measuring and comparing of alternatives to take place.

THE FACTORS THAT INFLUENCE PROCESSOR COSTS

Three main factors influence the cost of computer processors:

1. The current state of technology, which in effect defines what can be produced for a given cost.

2. The requirements and attitudes of users, which define the pattern of demand for processors.

3. The number, size and strategies of the competitors in the market, which define the products that are on offer and the prices that are asked for them.

The second and third factors in this list may be considered as "demand" and "supply" in the classical economic sense. The first factor, the state of technology, determines to a large extent the rate of change in the products that are bought and sold in the marketplace.

In most industries, technological change is a fairly slow and evolutionary factor that gradually shapes and improves the end product. In the computer industry, however, technology is usually viewed as being part of the product itself, and part of what is being sought and paid for. This means that technology is a powerful catalyst for change, and so it has a destabilising influence on the computer industry.

The cost of manufacturing a computer processor is closely related to the cost of its basic electronic components. Figure 1 shows how those costs have declined since 1955. The costs shown are those of a basic digital logic element (the flip-flop circuit), made by using three different manufacturing technologies. It is significant that each technology, although under pressure from a newer competitor, still managed to reduce its unit costs significantly over a period of a few years. It is also significant that the cost of each technology declined at a faster exponential rate than its predecessors: valves by 6% per annum, transistors by 9% per annum, and integrated circuits by 30% per annum.

The last curve on figure 1 shows the unit price (not the cost) of the Intel 8080 microprocessor chip. It was introduced in 1974, and from then until 1979 it declined in price by 60% per annum, or by almost a hundredfold over the five years. This device was exceptionally popular, and it belonged to the fastest-moving class of technological products of its time. Consequently, it is not truly representative of cost trends generally. Nevertheless, its price pattern is a typical, if somewhat accelerated, one.

The three manufacturing technologies correspond broadly to the so-called generations of computer processors that the market has witnessed. The first generation was based on vacuum-tube valve technology, the second on discrete transistors, and the third and subsequent generations (the exact number is a matter of opinion and dispute) on increasingly high-level integrated circuits.

Figure 1 Trend in the cost of electronics

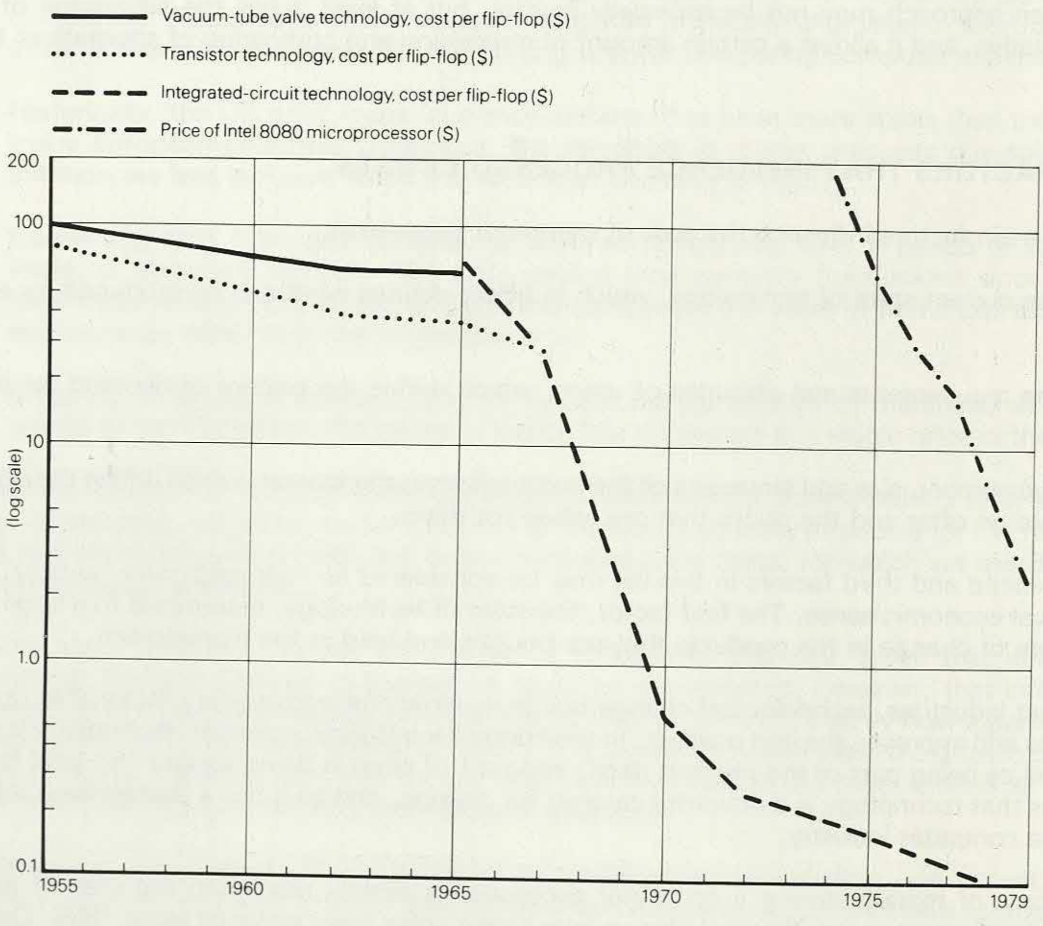

Each curve shows the historical cost trend of producing a flip-flop digital circuit using one of three technologies. Integrated circuits have produced by far the greatest cost reductions, and so they ushered in the era of cheap computer processors. The rightmost curve shows the price of the Intel 8080 microprocessor chip from the time it was introduced in 1974.

(Source: Phister)

However, well-established technologies do not disappear without a fight. No matter how persuasive are the economics of a new technology, history suggests that it will not automatically displace an existing technology. The computer marketplace has witnessed the continued commercial success of many anachronistic devices.

TRENDS IN PROCESSOR SPEEDS AND PRICES

Figure 2 gives a chronology of the trend in the cost of mainframe processor operations. The

Figure 2 Trend in the cost of processor operations

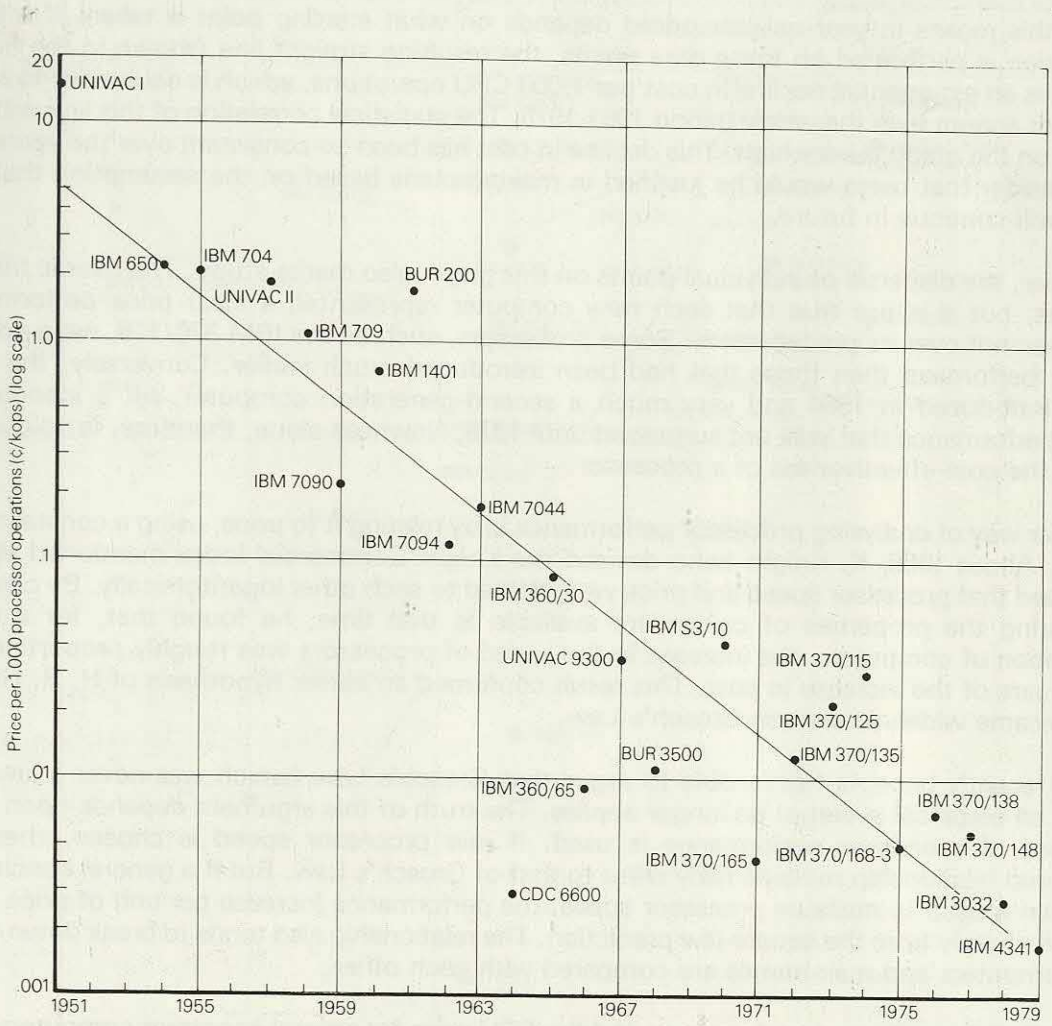

Each point shows the cost (in US ¢) of 1,000 processor operations on the computer indicated and in the year shown. The straight line is a least-squares fit to the points on the graph. Computers designated by points below the line therefore represented better-than-average value for money in their time.

(Sources: Auerbach and Phister)

scale is logarithmic and it displays a rather crude measure of computer power, namely the cost per thousand processor operations. This value is calculated from the purchase price of the processor and from a weighted measure (known as the Knight Commercial Index) that combines CPU execution speed, word size, and input/output speed. This value does not necessarily reflect the true cost of actually using the processor on genuine applications, but it is a consistent and composite measure which does not merely parrot the manufacturer's published "execution speed".

The computers included in this figure represent a spectrum of popular models spanning the years since the first commercial computers became available. A different selection of computers would yield slightly different results, but it is still clear that processor operations have declined in cost by three to four orders of magnitude in twenty-five years.

What this means in year-on-year prices depends on what starting point is taken. If a linear regression is performed on these data points, the resulting straight line (shown in the figure) indicates an exponential decline in cost per 1,000 CPU operations, which is equivalent to about 25% per annum over the whole period 1951-1979. The statistical correlation of this line with the points on the graph is very high. This decline in cost has been so consistent over the years that we consider that users would be justified in making plans based on the assumption that this trend will continue in future.

However, the dispersal of individual points on this graph also merits study. The overall trend is obvious, but it is not true that each new computer represented a clear price/performance improvement over its predecessors. Some processors, such as the IBM 370/115, were notably poorer performers than those that had been introduced much earlier. Conversely, the CDC 6600, introduced in 1964 and very much a second-generation computer, set a standard in price/performance that was not surpassed until 1978. Newness alone, therefore, is no guarantee of the cost-effectiveness of a processor.

Another way of analysing processor performance is by relating it to price, using a constant time factor. About 1968, K. Knight (who devised the Knight Commercial Index mentioned above) theorised that processor speed and price were related to each other logarithmically. By carefully measuring the properties of computers available at that time, he found that, for a given generation of computers, the increase in the speed of processors was roughly proportional to the square of the increase in cost. This result confirmed an earlier hypothesis of H. R. Grosch and became widely known as Grosch's Law.

It has recently become fashionable to argue that Grosch's Law (which was never a law, but rather an empirical average) no longer applies. The truth of this argument depends upon what measure of processor performance is used. If raw processor speed is chosen, then the measured relationship remains fairly close to that of Grosch's Law. But if a general benchmark program is used to measure processor speed, the performance increase per unit of price often falls off sharply from the square-law prediction. The relationship also tends to break down when minicomputers and mainframes are compared with each other.

Figure 3 shows values of performance and the CPU price for several important computers. The various generations of IBM processors are identified by the light symbols, and the non-IBM processors (not broken down by generation) are identified by the dark circles. In general, the further to the right and the lower down a point lies, the better is its price-performance.

One striking result is the fact that there was little if any improvement in price-performance between IBM's smaller 360 and 370 processor lines. The regression lines fitted to the points for the 360 and the 370 respectively are very close to each other, and they are almost exactly parallel. Although the 370 was introduced more than six years after the 360, its price/performance characteristics were only marginally better.

Figure 3 CPU price and performance for some popular computers and ranges

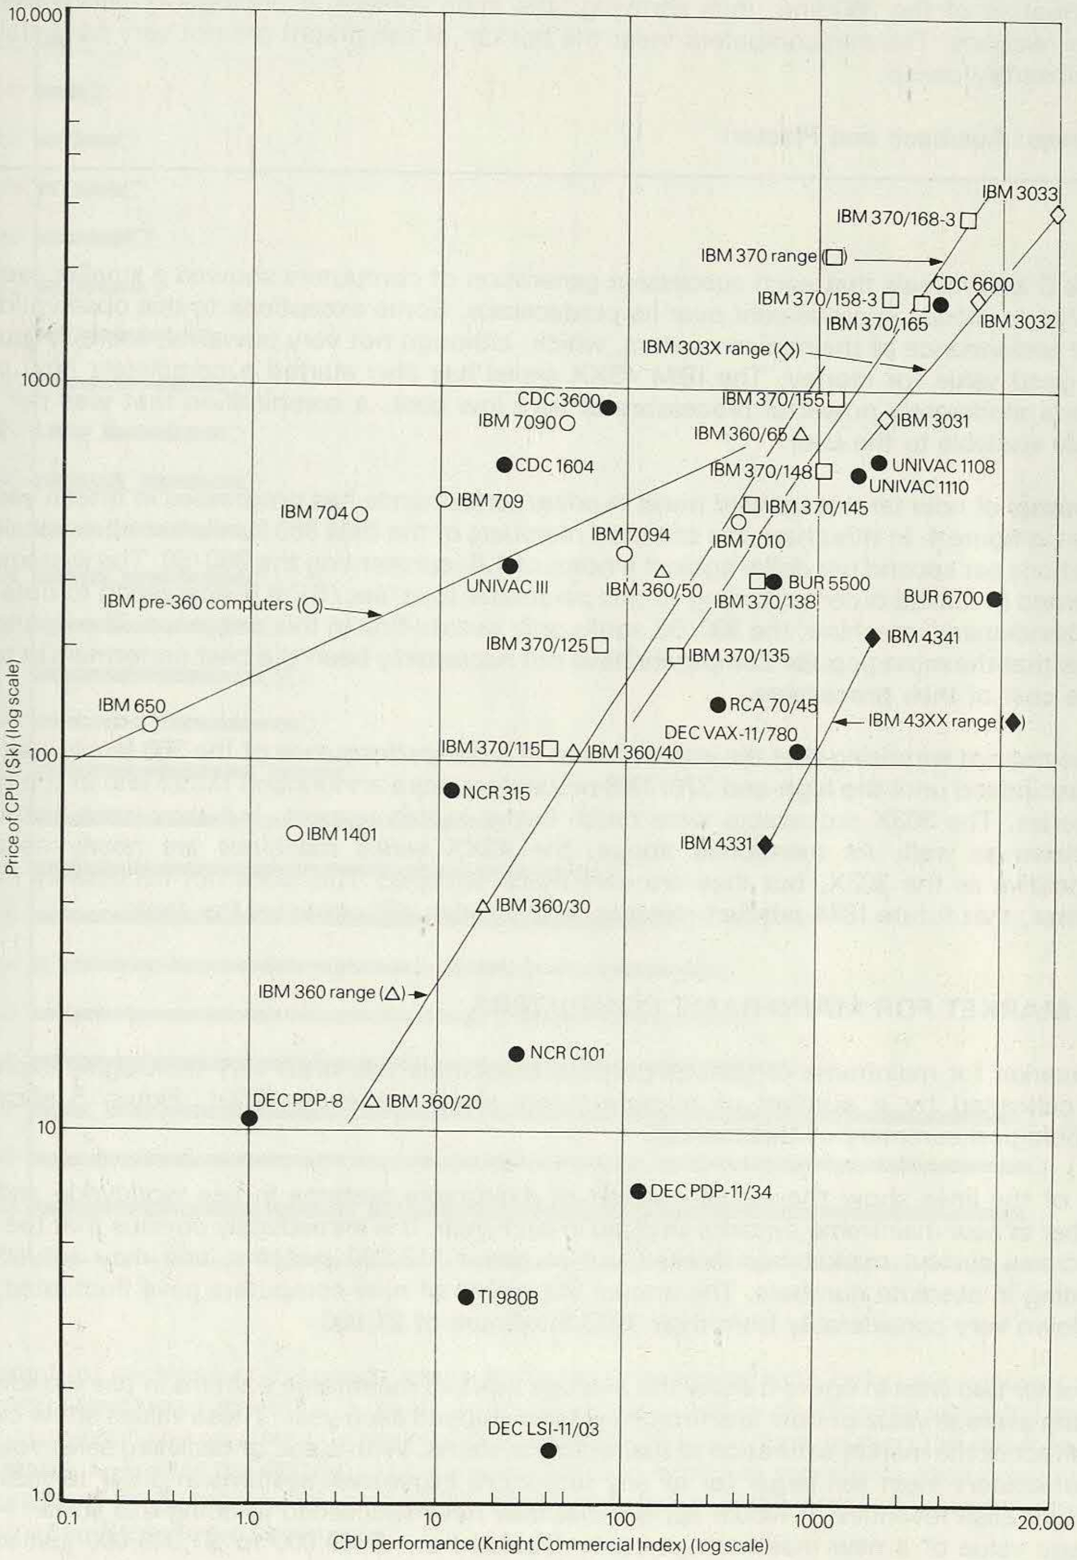

Each point represents the CPU performance index (x scale) and the price (y scale) of the computer indicated on the graph. Least-squares regression lines have been fitted to the pre-360, 360, 370, 303X, and 43XX ranges of IBM computers. Generally speaking, the further down and to the right a point lies, the better value it is. It is interesting that the 370 line is virtually a continuation of the 360 line, thus showing little improvement in price-performance over its earlier relations. The minicomputers (near the bottom of the graph) are not very powerful but are extremely cheap.

(Sources: Auerbach and Phister)

Figure 3 also reveals that each successive generation of computers showed a smaller average price/performance improvement over its predecessor. Some exceptions to this observation lie in the performance of the minicomputers, which, although not very powerful, seem to provide very good value for money. The IBM 43XX series has also started a completely new trend towards moderately powerful processors at very low cost, a combination that was not previously available to the user.

A measure of how far this upward trend in price/performance has progressed in fifteen years is shown in figure 4. In this chart, the standard members of the IBM 360 family have been scaled in operations per second per dollar against a norm of 1.0 representing the 360/50. The systems are presented in ranked order according to this parameter (ops/sec/$). It is interesting to note that the "benchmark" machine, the 360/50, ranks only twenty-first in this sequence. The figure also shows that the most popular computers have not necessarily been the best performers in terms of the cost of their processors.

It is somewhat surprising that the improvement in price/performance of the 360 family was very modest indeed until the high-end 370/1X8 processors were announced rather late in the life of the series. The 303X processors were much better in this respect, but they were also quite expensive as well. As mentioned above, the 43XX series machines are nearly as cost-competitive as the 303X, but they are very much cheaper. This does not necessarily mean, however, that future IBM product releases in this series will continue this trend.

THE MARKET FOR MAINFRAME COMPUTERS

The market for mainframe or general-purpose computers has been very thoroughly exploited and cultivated by a number of manufacturers since the early 1960s. Figure 5 shows a chronological summary of this market.

Two of the lines show the actual number of mainframe systems in use worldwide and the number of new mainframe systems shipped in each year. It is immediately obvious that the total mainframe system market has levelled out at about 113,000 systems, and may actually be declining in absolute numbers. The annual shipments of new computers have fluctuated, but are down very considerably from their 1973 maximum of 21,900.

The other two lines in figure 5 show the average value of mainframe systems in use worldwide, and the average value of new mainframe systems shipped each year. These values show clearly the effect of the market saturation in mainframe systems. With static or declining sales volume, manufacturers must sell larger (or at any rate more expensive) systems in order to maintain constant cash revenues. It would appear that they have succeeded in doing this so far — the average value of a new mainframe system doubled from $669,000 to $1,349,000 just in the period from 1974 to 1978.

One of the surprises caused by the 43XX series was that these computers, if they were sold only

Figure 4 Ranking of IBM 360 family computers by price-performance

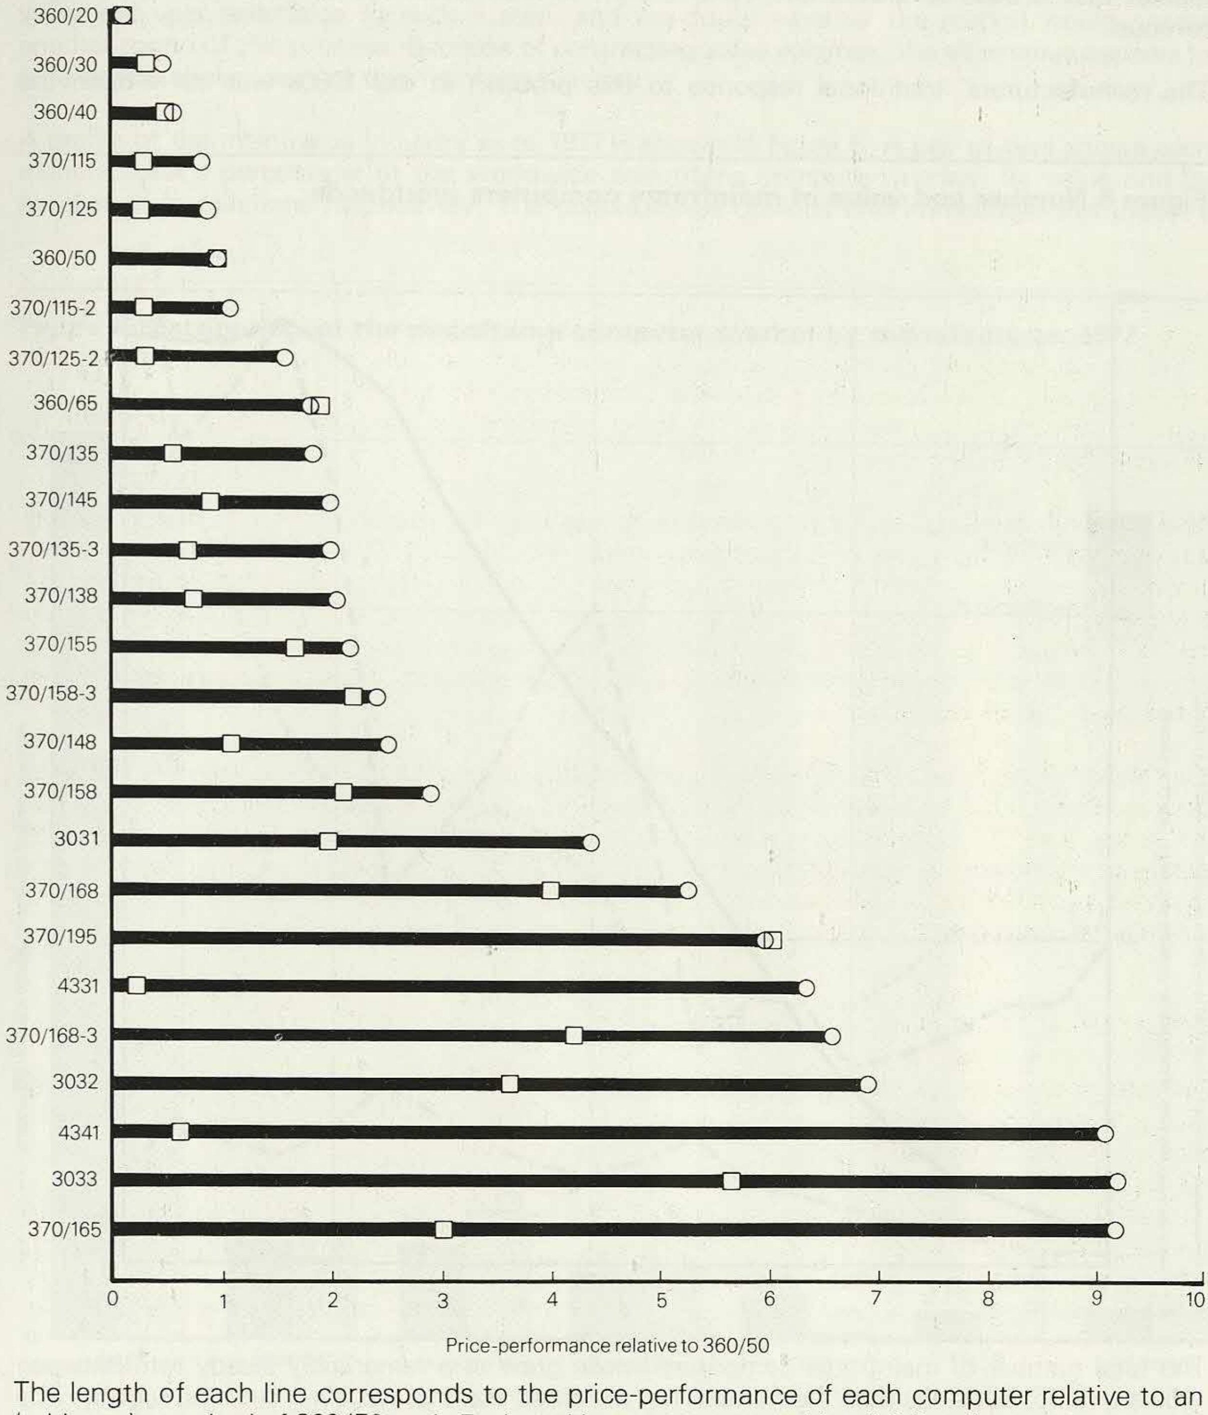

The length of each line corresponds to the price-performance of each computer relative to an (arbitrary) standard of 360/50 = 1. Each ranking was computed by dividing the weighted speed of the processor (in kops/sec) by the monthly rental (in $k), and then relating the result to the comparable figure for the 360/50. The square symbol in each line represents each computer's absolute cost (in $) relative to the 360/50. The order shown is the order of increasing price-performance (from top to bottom).

(Sources: Phister and Auerbach)

in traditional mainframe quantities, were not sufficiently expensive to generate IBM's traditional mainframe revenues. Consequently, IBM would either have to sell many more computers (into a market that is both saturated and contracting), or look elsewhere for additional sources of revenue.

The manufacturers' traditional response to this problem in the 1960s was to re-open the

Figure 5 Number and value of mainframe computers worldwide

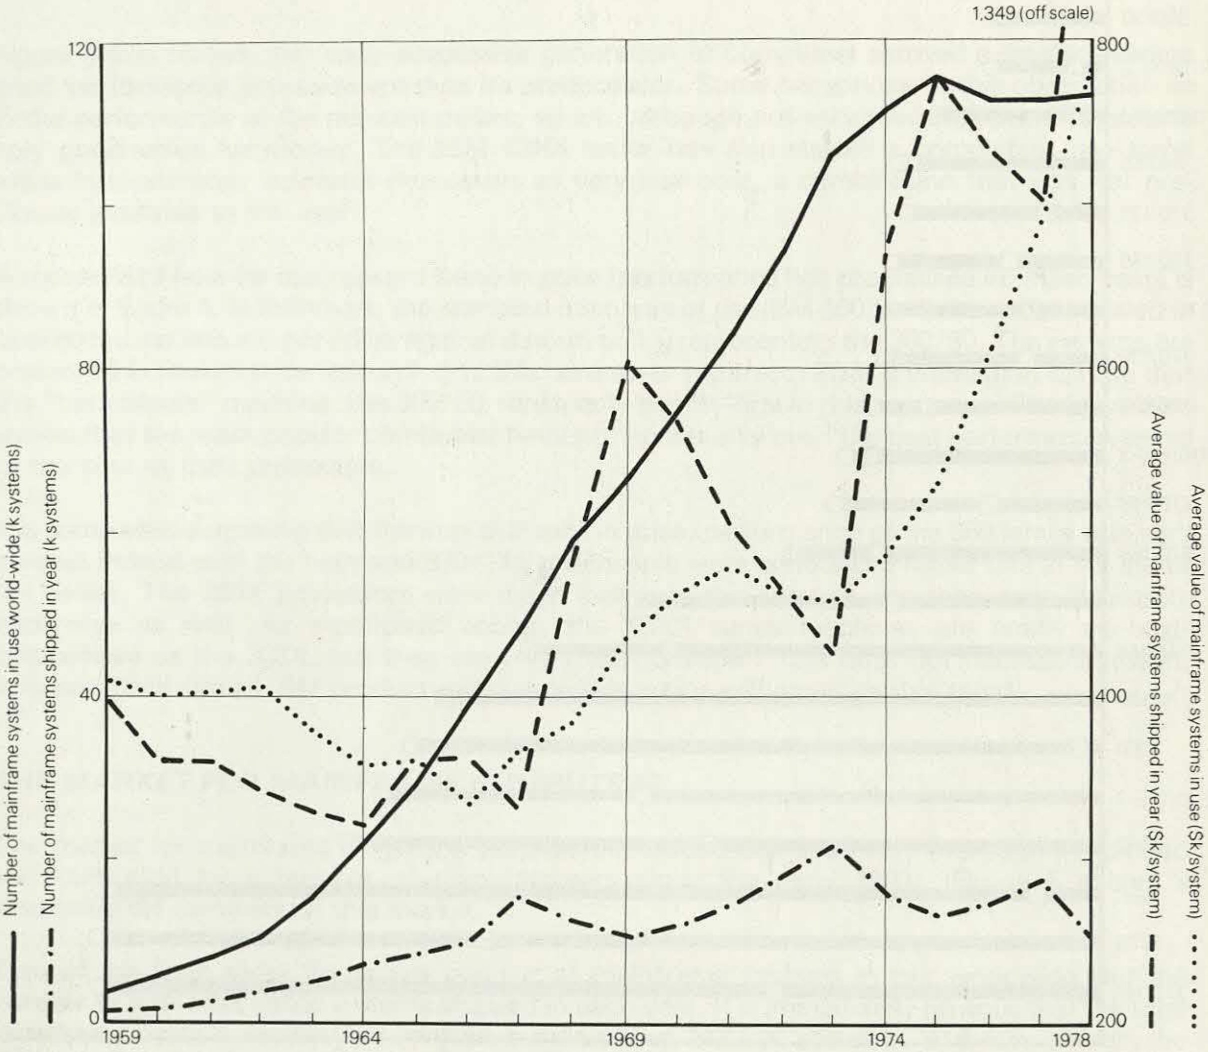

The total number of mainframe computers in use grew at a remarkably steady rate between 1959 and 1975, but since then the number has remained static or has even declined slightly. The shipments of new mainframes naturally reflects the introduction of new models and lines by manufacturers. The value of mainframes shipped has apparently been independent of the numbers shipped. It has varied enormously and in increasingly wide cycles. The average value of installed mainframes has increased very sharply since 1974, and it is off the top of the scale of this figure.

(Source: IDC)

saturated market by introducing a new, superior line of processors that was based on a new but incompatible architecture. Users then had to replace, rather than upgrade, their systems in order to utilise these new processors. In 1980, however, there appears to be widespread and vehement user resistance to such a step, and we doubt whether the market would accept another round of this process. Because of contracting sales volumes, the alternative appears to be a severe shakedown of the mainframe industry.

A profile of the mainframe industry as at 1977 is shown in figure 6. A pair of bars shows each manufacturer's percentage of the worldwide mainframe computer market, by value and by number of installations respectively. The conclusion is obvious and immediate: this industry

Figure 6 Distribution of the mainframe computer market by manufacturer, 1977

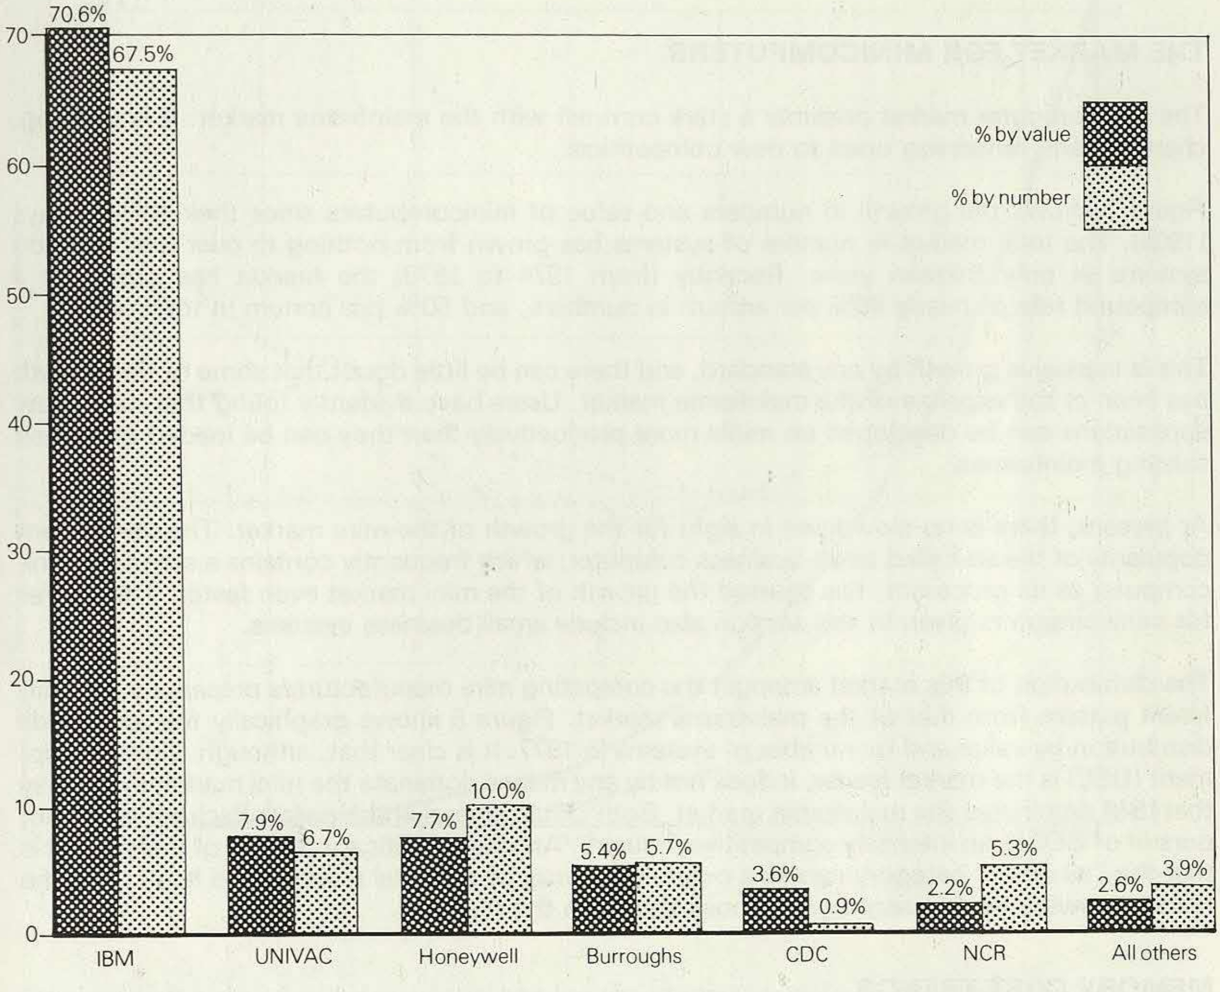

The height of each bar represents the percentage of the total worldwide mainframe computer market held by each manufacturer in 1977. The two bars in each position represent percentage by value and percentage by number respectively. The domination of the mainframe market by IBM is extremely evident.

(Source: Phister)

continues to be overwhelmingly dominated by IBM. An analysis of this industry breakdown over time (not shown in the figure) reveals that IBM's market share is in fact declining very slowly, as are the shares of most of its main competitors. The difference is being made up by a slow increase in the share of "all other", including all the plug-compatible manufacturers, as well as the Japanese companies, ICL, Siemens, etc.

If at first glance it appears unreasonable to bracket many familiar names in this category, it is important to remember that together they accounted for only 2.6% of the world total mainframe computer market by value in 1977, as compared with IBM's 70.6%.

No manufacturer has posed even a remotely serious challenge to IBM in this market since the mid-1960s. It is true that several computers have been introduced which offered much better cost-performance than IBM's lines (such as the Univac 1108, the CDC 6600, and the Burroughs B6700), but the manufacturers concerned could not maintain the pressure with successful follow-on products, and their attacks petered out.

THE MARKET FOR MINICOMPUTERS

The minicomputer market presents a stark contrast with the mainframe market. It is growing, changing and remaining open to new competition.

Figure 7 shows the growth in numbers and value of minicomputers since their earliest days (1965). The total market in number of systems has grown from nothing to over half a million systems in only thirteen years. Recently (from 1974 to 1978) the market has grown at a compound rate of nearly 40% per annum in numbers, and 50% per annum in total value.

This is explosive growth by any standard, and there can be little doubt that some of this growth has been at the expense of the mainframe market. Users have evidently found that many new applications can be developed on minis more productively than they can be loaded onto their existing mainframes.

At present, there is no slowdown in sight for the growth of the mini market. The very recent popularity of the so-called small-business computer, which frequently contains a standard minicomputer as its processor, has spurred the growth of the mini market even faster. The figures for minicomputers given in this section also include small business systems.

The distribution of this market amongst the competing mini manufacturers presents a very different pattern from that of the mainframe market. Figure 8 shows graphically the worldwide distribution by value and by number of systems in 1977. It is clear that, although Digital Equipment (DEC) is the market leader, it does not by any means dominate the mini market in the way that IBM dominates the mainframe market. Both Data General and Hewlett-Packard are in hot pursuit of DEC in an intensely competitive industry. Another significant feature of this market is that the "all other" category contains over one-quarter of the total market, and it is in fact the fastest-growing market segment of those shown in the figure.

MEMORY COST TRENDS

Computer memories have, in their history, undergone only one important conversion, but that conversion — from magnetic cores to integrated circuits in the early 1970s — was complete and far-reaching.

Figure 9 reveals clearly the trends leading up to this conversion, and indicates the crossover point. The vertical axis (log scale) shows the total manufacturing and testing costs of computer memory in dollars per thousand bits ($/kb). The three integrated circuit curves correspond to

1K, 4K, and 16K RAM chips. (A description of integrated circuit RAM technology is contained in Report No. 15, "Management Services and the Microprocessor".) By 1973, it was cheaper to make computer memory entirely from integrated circuit chips, and the manufacturers have never looked back on cores since then.

Figure 7 Number and value of minicomputers worldwide

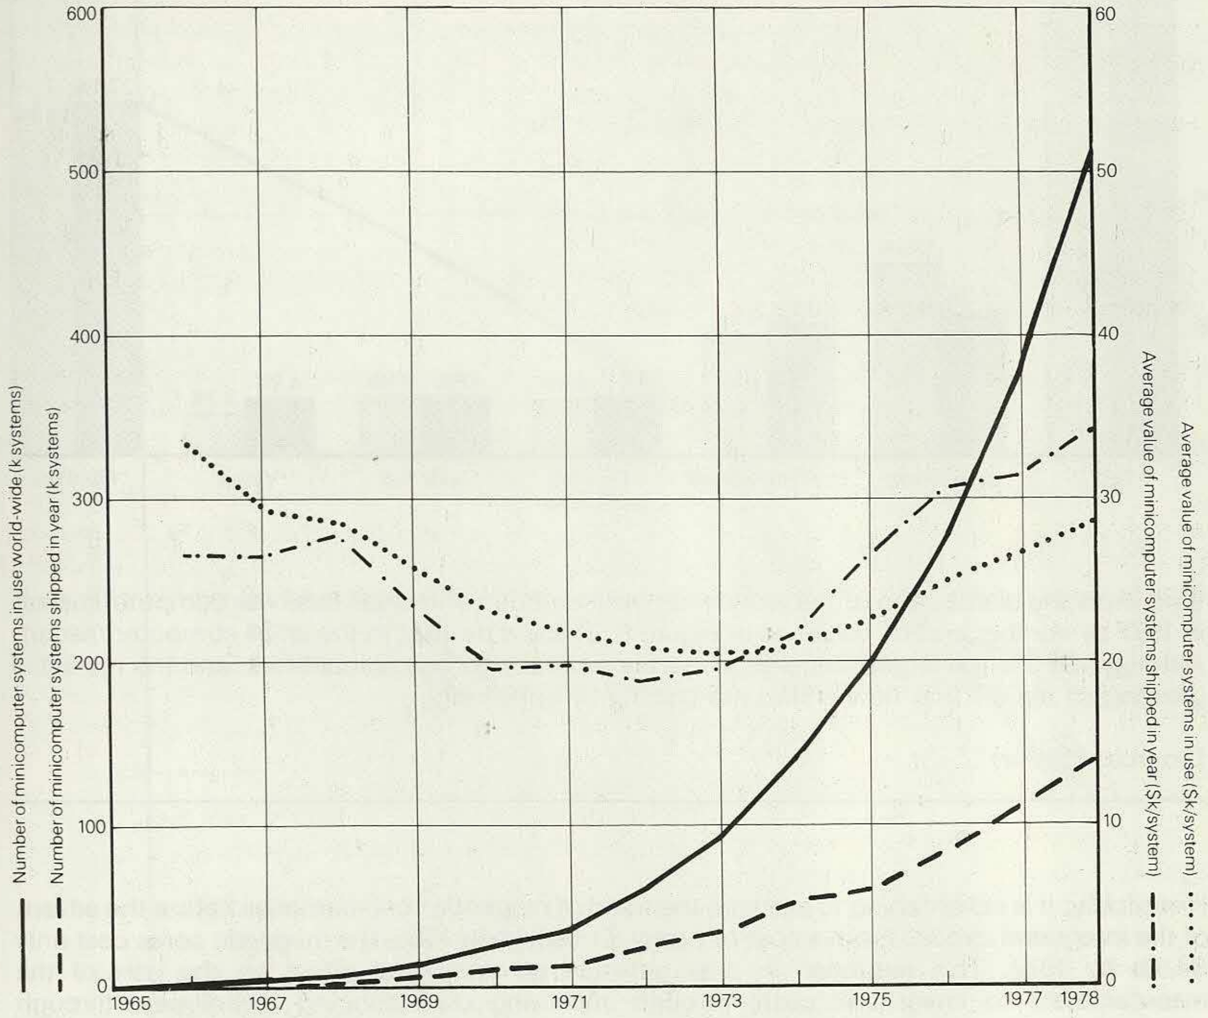

The installed base of minicomputers has increased exponentially in size since the mid-1960s. Moreover, the number of new shipments has increased every year, which means that the growth of the base is still accelerating. Prices of systems have generally declined until recently, when manufacturers sensed that the market was less price-sensitive than it previously was, and users began to configure mainframe systems from mini components. These figures include small business computers (SBCs).

(Source: IDC)

Figure 8 Distribution of minicomputer market by manufacturer, 1977

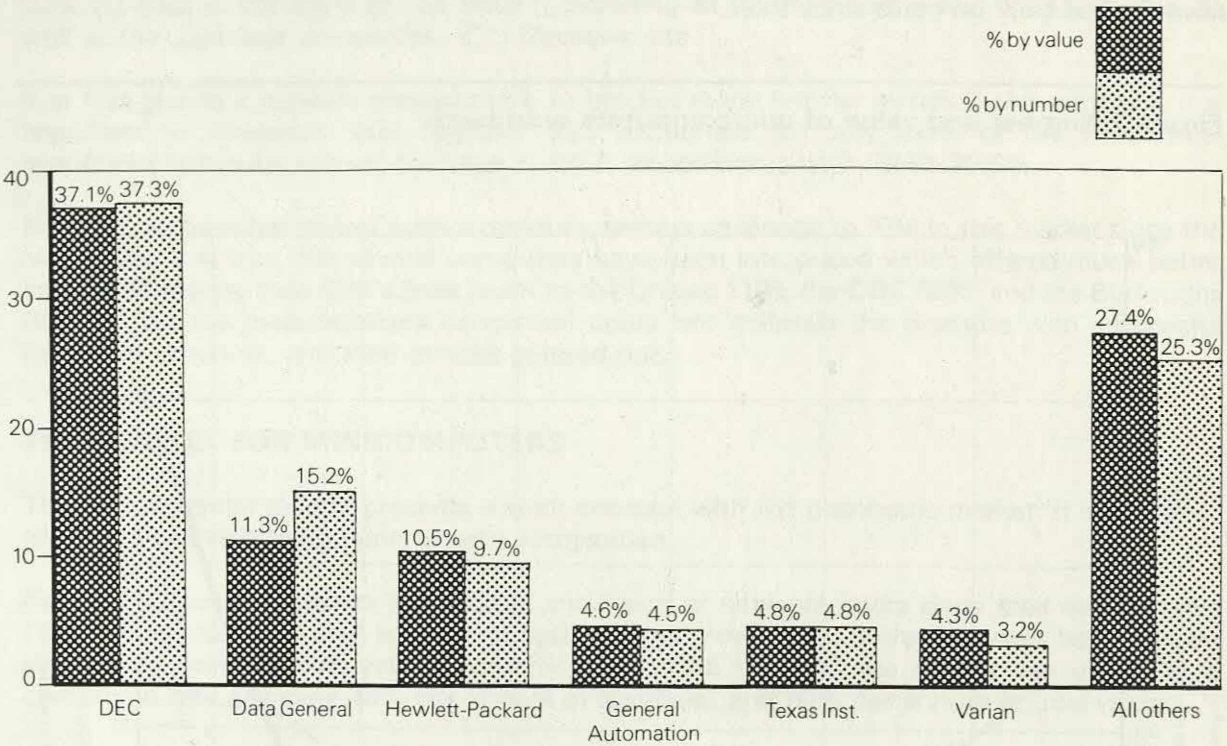

Bars show the distribution of the worldwide minicomputer and small business computer market in 1977 by number and by value, as in Figure 6. There is no IBM in the small computer market. Although DEC is the largest supplier, there are several vigorous competitors, and the market is growing so rapidly that new makes are constantly appearing.

(Source: Phister)

Historically, it is enlightening to examine the trend of magnetic core memories before the advent of the integrated circuit. From a cost of nearly $1,000/kb in 1955, the magnetic cores cost only $4/kb by 1972. This reduction in cost reflected an enormous effort on the part of the manufacturers to lower unit costs through improving core-stringing techniques, through mechanising assembly procedures, and even through farming out the most labour-intensive processes to Asian factories.

The result provides an excellent case history in the maturing of a technology in response to an increasing market demand for its product. Although integrated circuit memories proved to be inherently cheaper to make, the magnetic core, through sheer competitiveness, managed to stave off the inevitable integrated circuit victory until 1972.

The steady fall in the price of memory has had a disproportionate effect on computer systems, both in hardware and in the design of software and operating systems. Figure 10 indicates why this is true. The entries are in chronological order, and they show the price of memory in eighteen representative computers from the second generation onwards (the IBM 650 is included merely for reference).

Figure 9 Cost of manufacturing computer memories

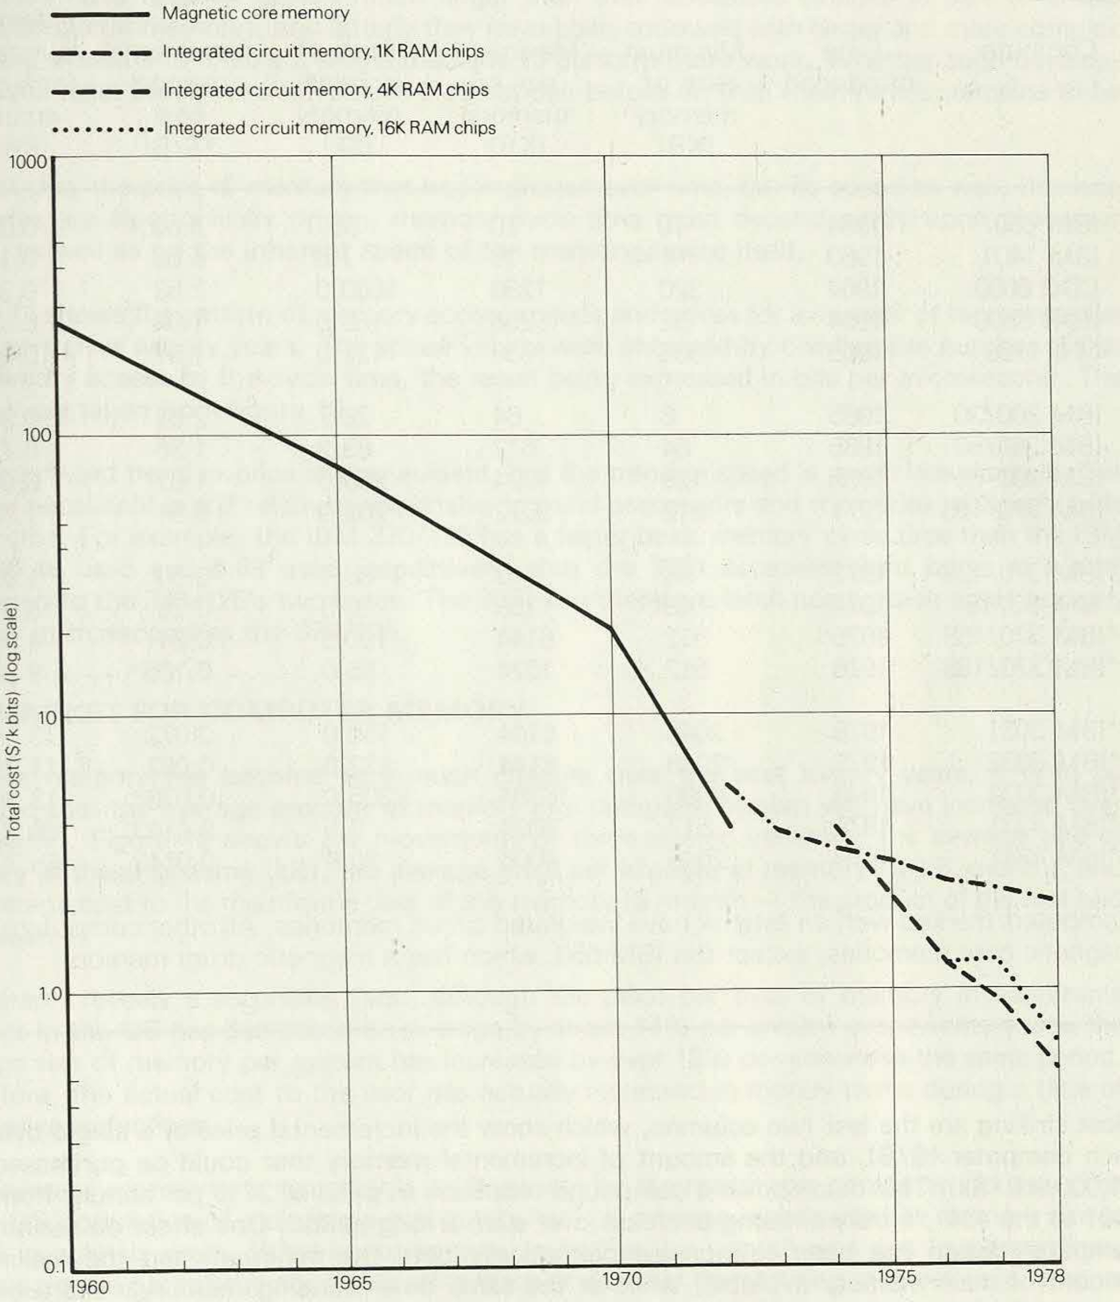

The curves show the total costs of manufacturing and testing computer memories, made with four different technologies, expressed in $/k bits. Whatever the inherent cost advantages of each technology, it takes several years for the unit cost to decline sufficiently to displace its predecessor. The decline in cost of core memory, prior to its being permanently displaced by integrated circuit memory, was spectacular. The trend of past curves suggests that the biggest improvements in integrated circuit memory are yet to come, provided that manufacturing problems can be smoothed out.

(Source: Phister)

Figure 10 Trends in memory capacity and price

Computer	Date introduced	Minimum size of memory (KB)	Maximum size of memory (KB)	Price of nominal memory ($k)	Incremental memory cost ($/B)	Memory purchased for a given amount (KB/$)
IBM 650	1954	10	20	35.0	3.50	0.28
IBM 1401	1960	1.4	16	7.8	5.58	0.18
CDC 6600	1964	320	1280	1660.0	2.53	0.39
BUR 5500	1964	32	256	180.6	1.84	0.53
UNI 1108	1965	384	1536	421.0	1.07	0.91
IBM 360/30	1965	8	64	20.7	2.52	0.39
IBM 360/50	1965	64	512	89.2	1.36	0.72
IBM 360/65	1965	128	1024	200.0	1.52	0.64
IBM 370/165	1971	512	3072	269.0	0.51	1.92
*IBM 370/135	1972	96	512	88.1	0.90	1.09
*IBM S3/15	1974	48	256	12.5	0.25	3.91
*IBM 370/158	1975	512	6144	126.5	0.24	4.07
*IBM 370/138	1976	512	1024	55.0	0.105	9.30
*IBM 3031	1978	2048	6144	150.0	0.072	13.56
*IBM 3032	1978	2048	6144	172.0	0.082	11.91
*IBM 3033	1978	4096	16384	305.0	0.073	13.38
*IBM 4331	1979	512	1024	7.5	0.0143	68.29
*IBM 4341	1979	2048	4096	30.0	0.0143	68.29

Computers marked with an asterisk have integrated circuit memories. All other computers have magnetic core memories, except the IBM 650, which has a magnetic drum memory.

Most striking are the last two columns, which show the incremental price of a single byte for each computer ($/B), and the amount of incremental memory that could be purchased for $1,000 (kB/$k). This data shows a compound reduction in price of 27% per annum from the 1401 to the 4341, a truly amazing decrease over such a long period. One effect on mainframe computer design has been to increase dramatically both the minimum and the maximum amount of main memory available, while at the same time reducing memory's share of the overall system price.

However, a great deal of software had been written over the years for the purpose of alleviating the shortage of the expensive memory in which programs had to reside. These ranged from simple overlays through swapping and paging to complete virtual-memory systems. Although memory is now cheap and plentiful ($15,000/megabyte in the IBM 4341), the legacy of cramped, expensive memory survives today in the form of inappropriate software. Our conclusion is that such software — which is running today on many, if not most, mainframe computers — is oriented towards solving the wrong problem. As a result, users are paying to remove a bottleneck that need no longer exist, and to do so they must support large and complex operating systems that may actually reduce the effectiveness of their computers.

One of the attractions of the minicomputer has been the relative absence of such operating systems. Originally, this software was not necessary, because the minicomputer did not attempt to perform more than one simple task at a time. Ironically, however, some minicomputers have recently grown much larger than their ancestors (thanks in part to cheap integrated circuit memory), and latterly they have been endowed with larger and more complex operating systems to enable a single machine to perform more work. Whether such developments will repel users, who have seen it all happen before on their mainframes, remains to be seen.

It is not only the price of memory that has improved over time, but its speed as well. Because memories are electronically driven, memory cycle time must depend partly upon processor speed, as well as on the inherent speed of the memory device itself.

Figure 11 shows the pattern of memory access speeds and prices for a number of representative computers over twenty years. The speed values were obtained by dividing the number of bits per memory access by the cycle time, the result being expressed in bits per microsecond. The prices were taken from figure 10.

The downward trend in price is very evident, but the trend in speed is much less marked. This may be because it is still relatively expensive to build processors and memories with very wide data paths. For example, the IBM 370/125 has a faster basic memory cycle time than the IBM 3031 (0.48 usec and 0.69 usec respectively), but the 3031 accesses eight bytes at a time compared to the 370/125's two bytes. The 3031 can therefore fetch nearly three times as many bits per microsecond as the 370/125.

THE MARKET FOR COMPUTER MEMORY

Because memory has become very much cheaper over the past twenty years, it is to be expected that the average amount of memory in a computer system will have increased over that period. Figure 12 shows the movements of three related variables: the average size of memory in these systems (kB), the average price per kilobyte of memory ($/kB/month), and the average cost to the mainframe user of this memory ($/month — the product of the first two variables).

This graph reveals a surprising fact: although the price per byte of memory in mainframe systems in the US has decreased on average by about 14% per annum over twenty years, the average size of memory per system has increased by over 18% per annum in the same period. Therefore, the actual cost to the user has actually increased in money terms during a time of rapid price reductions.

This increase represents a remarkable achievement by the mainframe computer industry. The entire US population of mainframe computers has, on average, quadrupled its memory size to 426kB in only six years. What is not known is whether this growth was due to extraordinary success on the part of the computer sales force or whether the growing memory requirements of operating systems and applications absorbed the increase. What can be said from the statistical evidence is that users today routinely employ amounts of memory that were unheard of a decade ago.

EXPENDITURE ON PROCESSORS AND MEMORY

Users' actual expenditure on processors and memory over the past two decades is summarised in figure 13. The height of each bar in the chart is composed of two parts: the average US mainframe user's expenditure on *processor* and on *memory* for each year. Annual expenditure was estimated by dividing purchase price by 44 to obtain a notional monthly rental, and then multiplying by 12 to obtain an annual cost.

The fluctuations in users' expenditure on processors are quite pronounced, varying from a low of $15.2k in 1966 to $54.6k in 1978 on a steeply-rising trend. The 1978 figure was double that for 1975.

Expenditure on memory has been notably consistent, despite huge changes in the price per byte. Indeed, the expenditures in 1977 and 1971 were identical, although massive amounts of new memory were installed over that period.

Figure 11 Trends in memory cost and performance

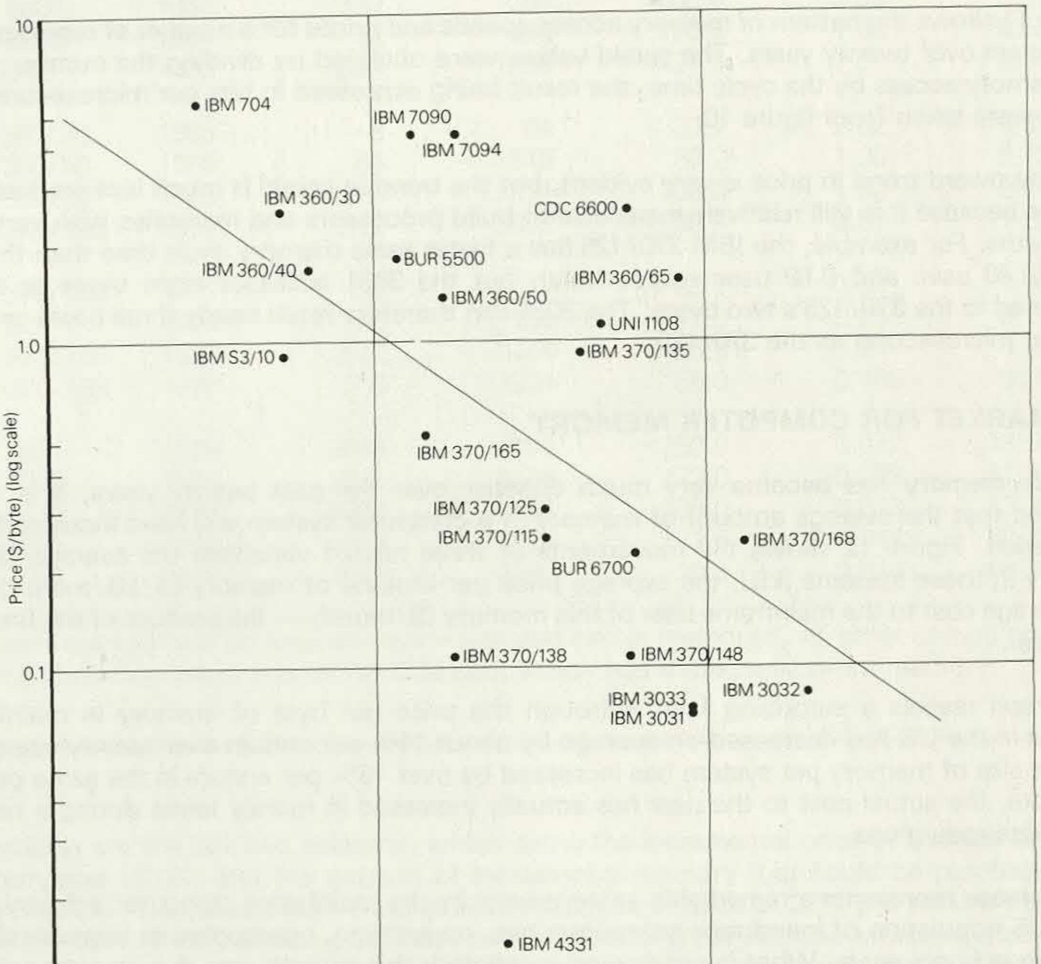

Each point on the graph shows the speed and the price of the memory associated with a particular computer. The speed is a composite figure of bits per microsecond, derived by dividing the number of bits per memory access by the cycle time. Price is expressed in $/byte. The regression line fitted to these points shows the close relationship between these variables over the period 1955 to 1979 represented by these computers. Both scales are logarithmic.

(Source: Phister)

Figure 12 Size and price of average user memory configuration

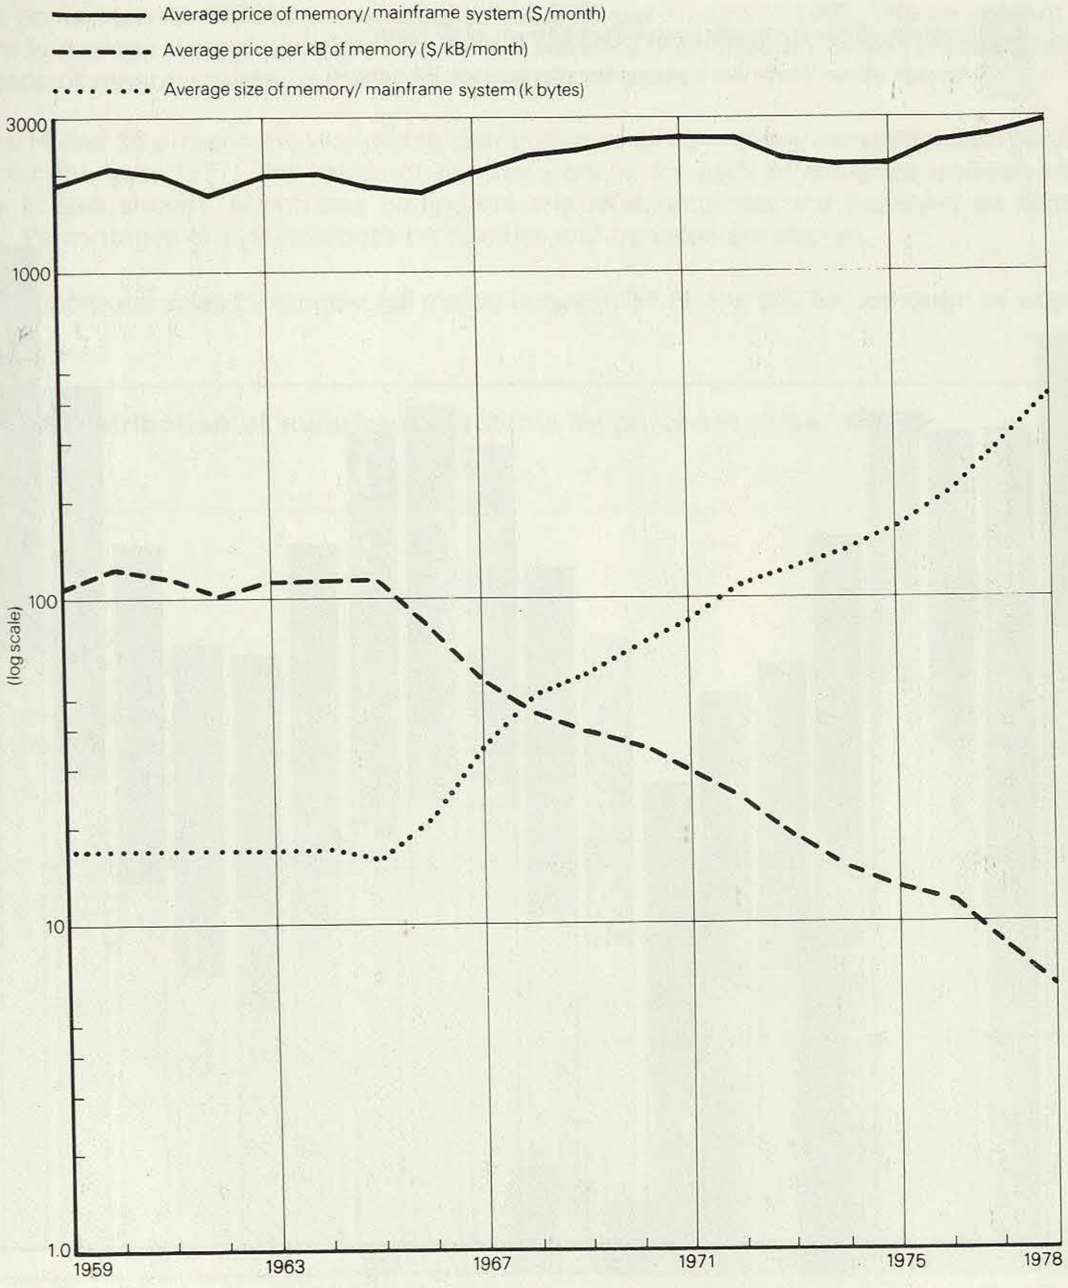

The curves show how the average price of memory per month has declined, but has been more than offset by the increase in the average amount of memory per mainframe computer system. The top curve, the product of the other two curves, shows that the net total user expenditure on memory has in fact increased in absolute terms, despite the dramatic fall in unit price. These figures are for mainframe computer installations in the US.

(Source: Phister)

Figure 13 Average user expenditure on processors and memory

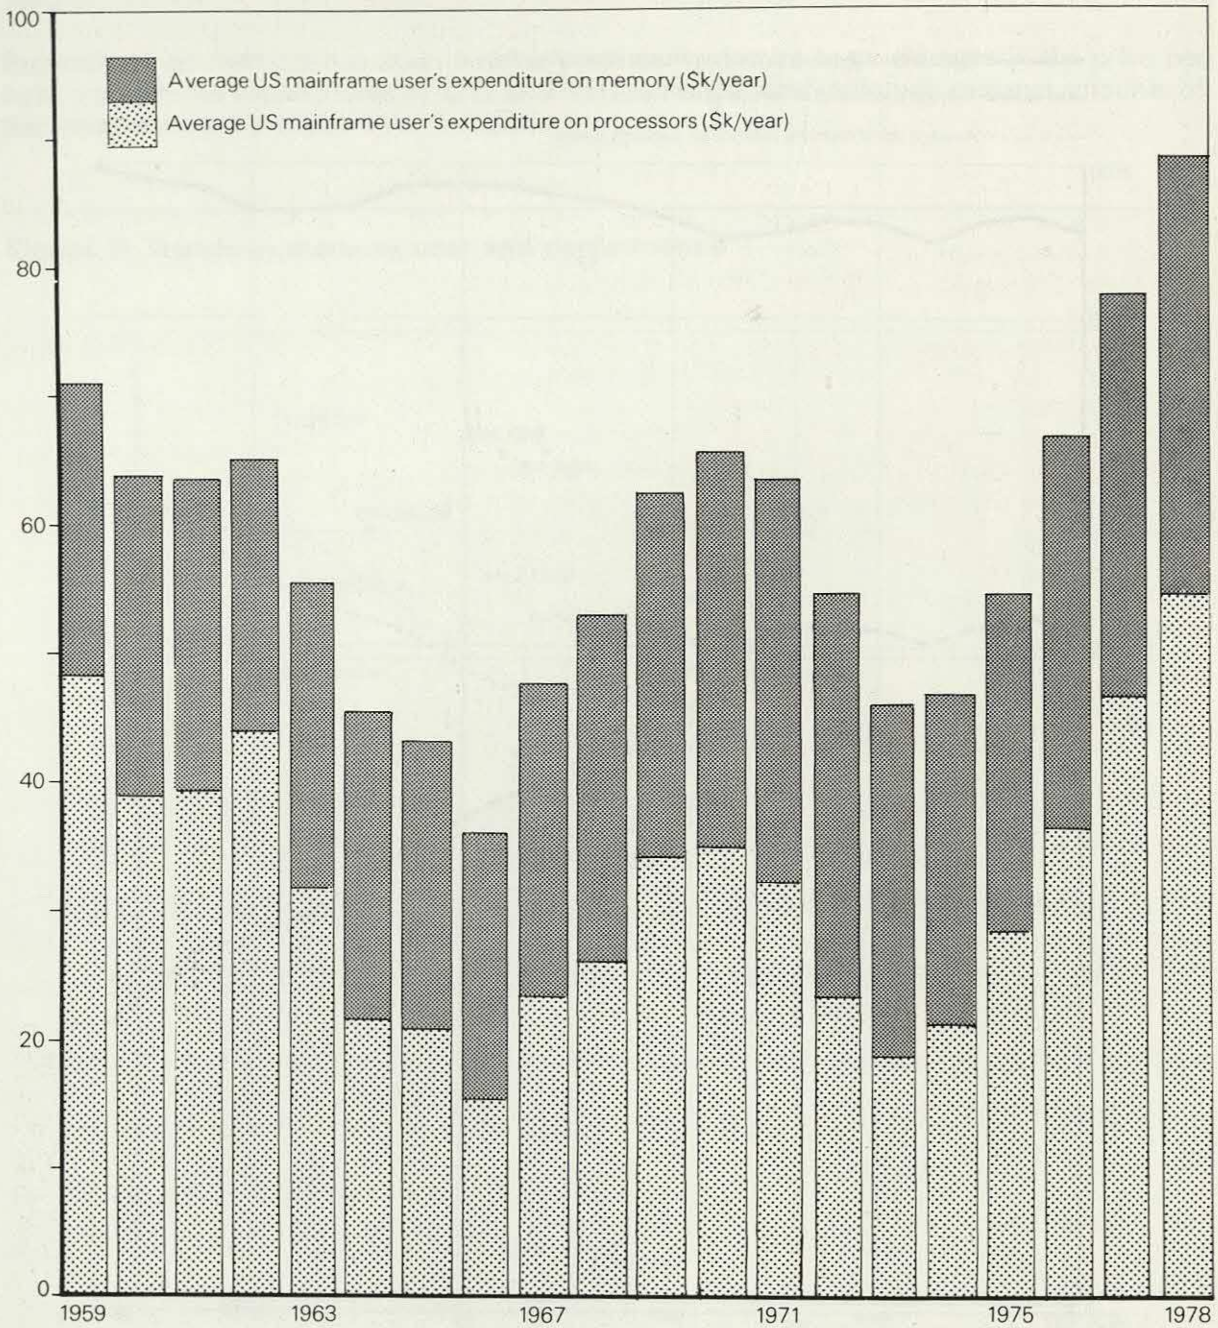

The figure shows what mainframe users in the US actually spent, on average, on processors and memory in the years given. The height of the lower part of each bar indicates expenditure on processor, while the remainder above it indicates expenditure on memory. The height of each full bar therefore represents the total of processor and memory. Expenditure on processors has gone through several major cycles since 1959, and it is still on a very strong upward surge in 1978. Spending on memory has fluctuated less, but it is still higher now than ever before, despite dramatic reductions in the cost of computer memory.

(Source: Phister)

Total expenditure on processor and memory (excluding the two most recent years 1977 and 1978) has varied by less than a factor of two over the entire period. This is remarkable in view of the changes to the computing industry that occurred during this time, and in view also of the massive growth in the number and variety of systems installed. Nevertheless, the figure of \$66.8k per annum for 1976 is nearly equal to the \$65k per annum of 1962. The very recent leap upward in average system price appears to be a separate phenomenon, which probably reflects the needs of manufacturers, rather than the needs of users.

Figures 14 and 15 present the worldwide distribution of all new computer systems by purchase price, for the year 1977. The range of purchase prices for each of the price intervals of the figures is also shown. Mainframe computers and minicomputers are displayed on separate charts. Percentages of systems both by number and by value are shown.

The minicomputer sales by number fall mainly between \$7.8k and \$62.5k, although by value the

Figure 14 Distribution of mainframe systems by purchase price, 1977

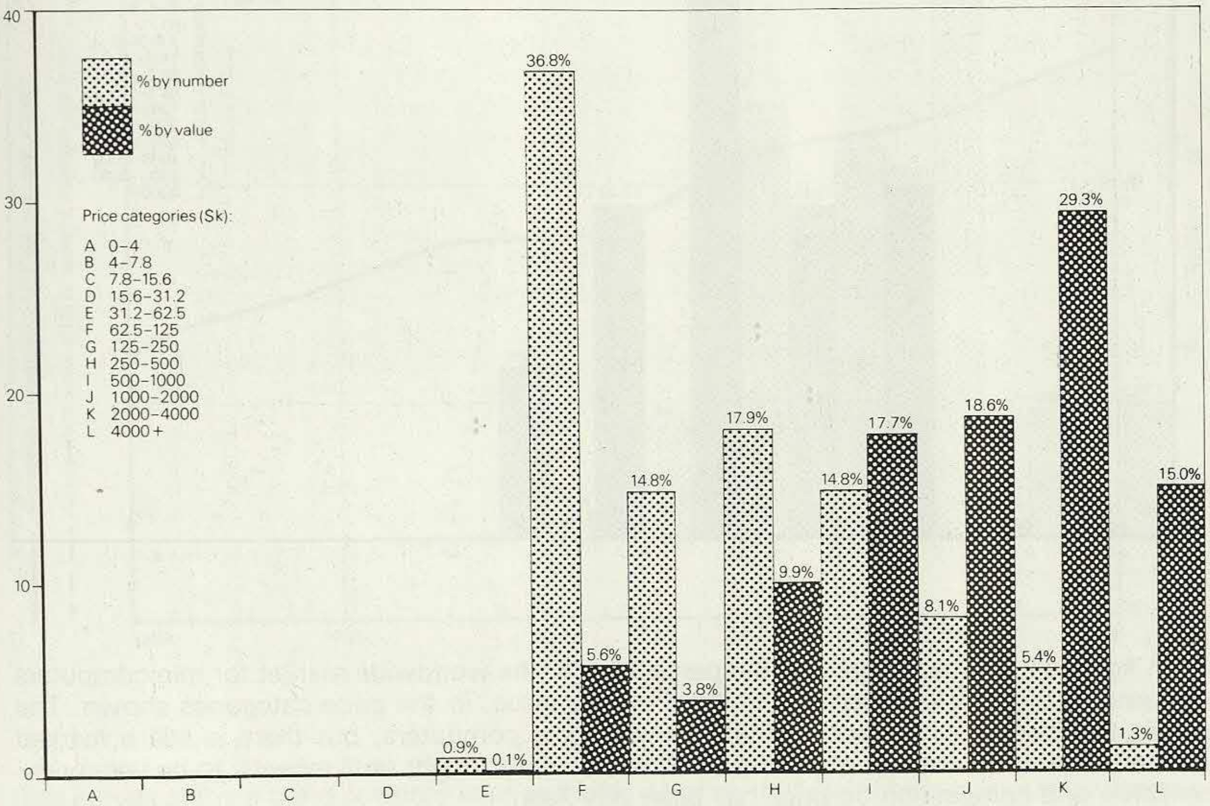

Each pair of bars on the chart shows the percentage of the worldwide mainframe computer population that falls within the price category indicated. The left bar in each pair is the percentage by total number; the right bar is the percentage by total value. The market is dominated by purchases of cheaper computers, but machines costing over \$2m nevertheless account for nearly half the total by value.

(Source: IDC)

larger systems predominate. With mainframe computers, this tendency is much more pronounced. Although 52% of mainframe systems by number cost less than $125k each, 44% by value are accounted for by systems costing over $2M each. Clearly, the market for very large systems still ranks very high in importance amongst the mainframe manufacturers, even though the sales of such machines are few in number (less than 7% of all mainframe sales).

Figure 15 Distribution of minicomputers and small business computers by purchase price, 1977

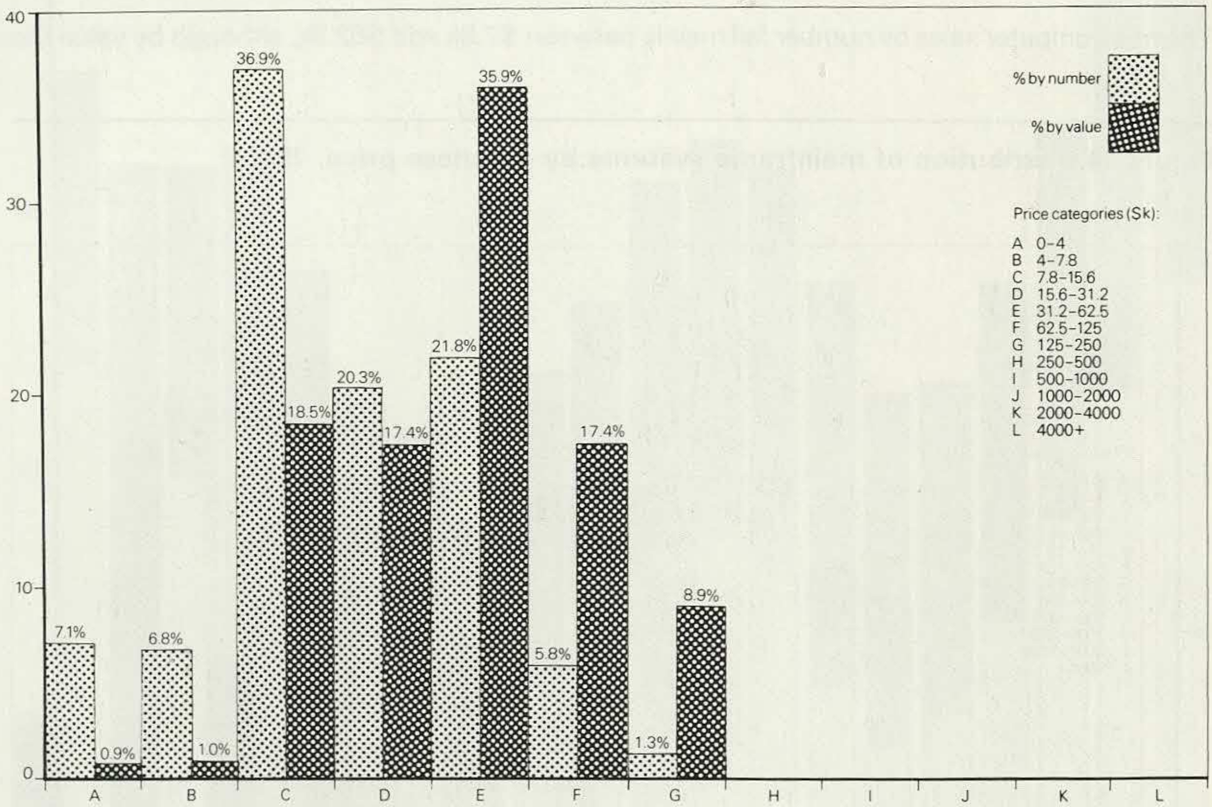

As in figure 14, the bars represent the percentage of the worldwide market for minicomputers and small business computers, by number and by value, in the price categories shown. The distribution is not so extreme as that for mainframe computers, but there is still a marked imbalance between number and value. This market was thought until recently to be very price-sensitive, and competition on price has been very keen.

(Source: Phister)

Finally, figure 16 shows the trend in the average service life of all mainframe systems shipped in the US to date, and also that of systems being retired in each year. In the long term, these two lines could be expected to converge.

Also shown is a more volatile measure, each year's system retirements expressed as a

Figure 16 Mainframe system life and system retirements

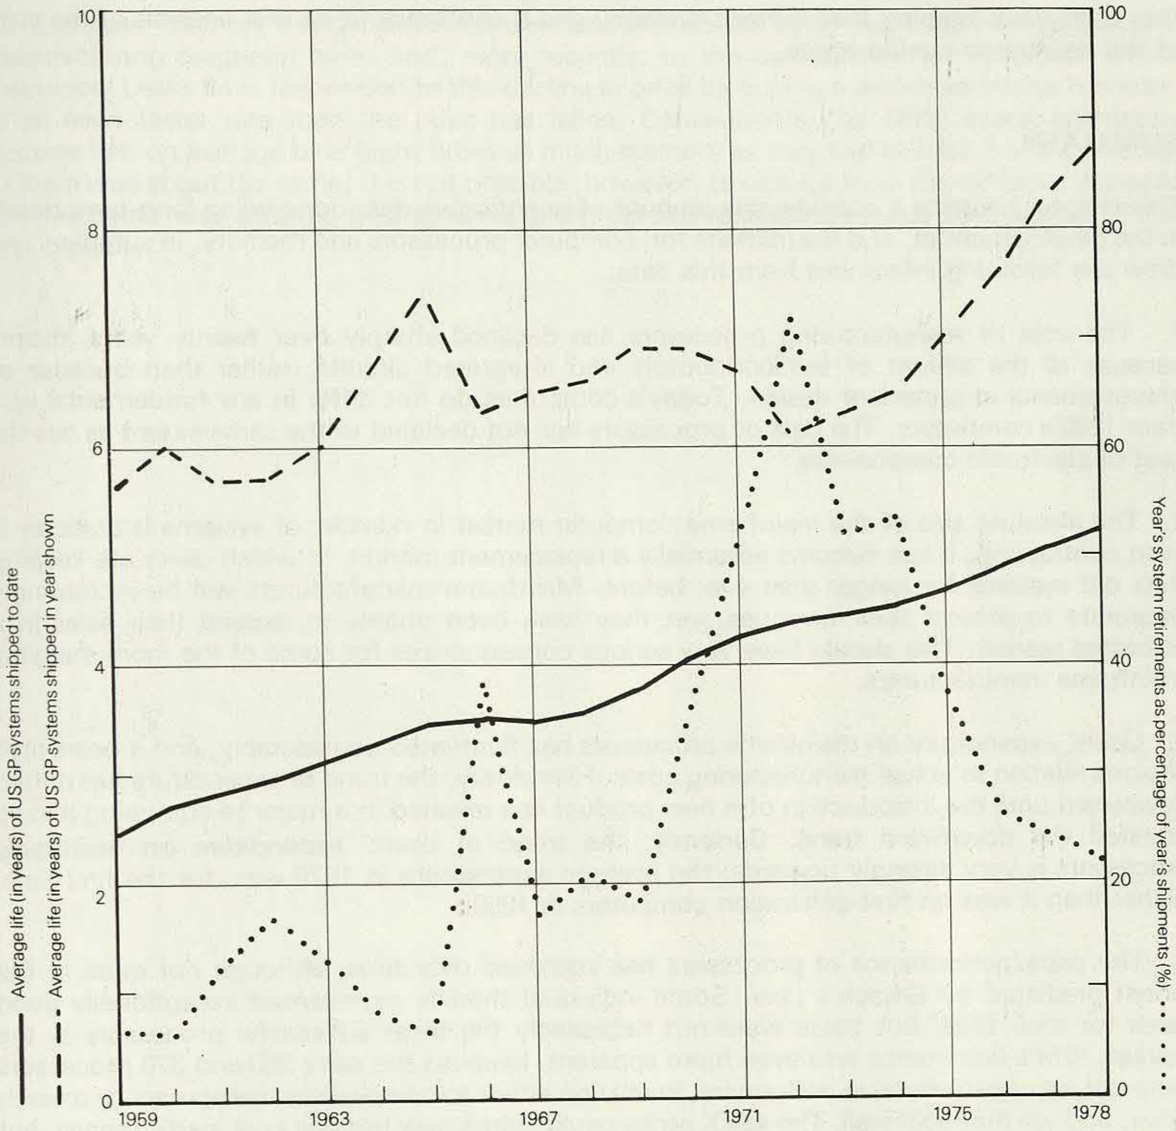

The curves show a trend towards later and later retirement of mainframe computer systems in the US. If users are keeping their machines longer, and if the market for mainframe systems is not growing, then the prospect for mainframe computer manufacturers seems bleak. The third curve shows the retirements of mainframe systems expressed as a percentage of the shipments of new systems in the same year. This very volatile measure captures the impact of the introduction of new computer lines. Users traded in their old systems by the thousands when the 360 and 370 series were introduced. The market now seems past due for another such cycle, if the past trend is to be repeated.

(Source: IDC)

percentage of the same year's new shipments. This in effect measures how rapidly users are abandoning their old systems for new ones. This curve naturally shows pronounced peaks at those times when users were rapidly trading up for new systems, namely at the beginning of the second generation (1961), after the launch of the 360 series (1964), and after the launch of the 370 (1970). Since 1972, this trend has moved steeply and steadily downwards. This indicates that users were keeping their current systems, and it confirms our earliest analysis of the state of the mainframe market today.

SUMMARY

This chapter contains a considerable amount of quantitative data concerning long-term trends in the development of, and the markets for, computer processors and memory. In summary, we draw the following inferences from this data:

1. The cost of manufacturing processors has declined sharply over twenty years, mainly because of the advent of semiconductors and integrated circuits, rather than because of improvements in processor design. Today's computers do not differ in any fundamental way from 1960's computers. The cost of processors has not declined to the same extent as has the cost of electronic components.

2. The absolute size of the mainframe computer market in number of systems is static or is even contracting. It has become essentially a replacement market, in which users are keeping their old systems for longer than ever before. Mainframe manufacturers will be increasingly desperate to protect their revenues, yet they have been unable to expand their sales in a saturated market. This should have very serious consequences for some of the more marginal mainframe manufacturers.

3. Users' expenditure on mainframe processors has fluctuated considerably, and it bears little obvious relation to actual manufacturing costs. Historically, the trend of expenditure has drifted downward until the introduction of a new product line resulted in a major re-equipping and so arrested the downward trend. Currently, the trend of users' expenditure on mainframe processors is very strongly upwards: the average expenditure in 1978 was, for the first time, higher than it was on first-generation computers in 1959.

4. The price/performance of processors has improved over time, although not quite to the extent predicted by Grosch's Law. Some individual models represented exceptionally good value for their time, but these were not necessarily the most successful processors in the market. IBM's dominance was even more apparent, because the early 360 and 370 processors were not very cost-effective performers, in relation either to competitors' computers or to each other, and yet they sold well. The 43XX series could start a new trend in cost/performance, but whether this will happen probably depends more on the decisions IBM takes than on events in the marketplace.

5. The minicomputer market has taken off at a terrific rate, and no slowdown is yet in evidence. There are many manufacturers' products in this market, and the competition is intense. Until recently, users were thought to be highly price-sensitive, but there is evidence that they are now less so, as minis are used for more and more commerical DP tasks. It seems that some of the mini's spectacular growth in sales has been at the expense of the mainframe market.

6. Despite sustained competition from several other companies, IBM still firmly dominates the mainframe market, with 70% of the market by value. This dominance is even more pronounced if one considers that several of the competitors are producing copies of IBM computers. The minicomputer market, however, is not so concentrated. The market leader, DEC, has only 37%

of the total, and this share is declining. Many other companies are active in this expanding market, and the cost of entry is relatively low. Because minicomputers are used for so many kinds of application, there is scope for many different manufacturers with different marketing strategies.

7. Computer memory has declined exponentially in price per byte, thanks to improvements in manufacturing magnetic cores and, more recently, to the development of integrated circuit memories. Users have responded to this decline in price by buying massive additional memory, at an even faster rate than the price has fallen. Consequently, by 1978, users' mainframe systems had on average over eight times as much memory as they had in 1968, but the net cost to them was about the same. It is not possible, however, to deduce from the evidence available whether users have adapted their systems and programs to make the most effective use of this vast capacity.

CHAPTER 3

AUXILIARY STORAGE DEVICES

By auxiliary storage devices we mean mainly magnetic discs and tapes, those stalwarts of the computer industry that almost invariably accompany the installation of a new computer. Their importance has grown considerably since the earliest days of computing, when the processor was all-important and peripherals were purely a means to an end. Now the emphasis has changed towards data-driven systems, database management is in vogue, and auxiliary storage devices are much nearer the centre of activity.

TRENDS IN DISC STORAGE

In discussing magnetic disc units, we confine the subject to moving-head disc files, because these are by far the most popular units in use today. Other devices, employing a drum or head-per-track approach, enjoyed a few years of popularity during the 1960s, but they faded away rapidly as their costs grew with increased capacity. They are now practically obsolete.

The earliest disc units were complex and slow, because a single head was employed to service every location within the multi-surface disc structure. The first practical disc device to reach the market was the IBM 1301 in 1961. This model provided multiple heads to accommodate multiple recording surfaces, and so it simplified the access procedure and reduced the average access time five-fold over the earlier devices. At the same time, the recording density was increased, and this helped to offset the cost of the new mechanism. The cost per byte remained about the same, whilst speed and capacity were improved.

An additional innovation followed in 1963 in the form of the IBM 1311. This model featured a removable disc "pack", which was analogous to a reel of magnetic tape in that it offered off-line storage as well as rapid on-line accessibility on the same recording medium.

Subsequent developments in disc products have followed similar lines by providing higher recording density, smaller physical storage media and marginally faster access mechanisms. The results can be seen in figure 17, which shows the trends in cost per byte, spindle capacity, and effective data transfer rates for several IBM disc products over the years.

The only other significant development has been the reversion to fixed media (that is, non-exchangeable disc units). The removable disc packs always incurred a stiff price penalty compared with the fixed units and, because of head positioning problems, were often less reliable.

Although the improvements in moving-head disc price-performance have been evolutionary rather than startling, they have also been both continuous and cumulative. As a result, disc units have improved over a period of twenty years by two orders of magnitude in price per byte, and by nearly as much in both capacity per spindle and data transfer rate. At the same time, the price per spindle has declined steadily. This decline has given rise to a buyer's market, in which improved performance and lower absolute prices progressed together.

Figure 17 Trends in disc price, capacity and speed

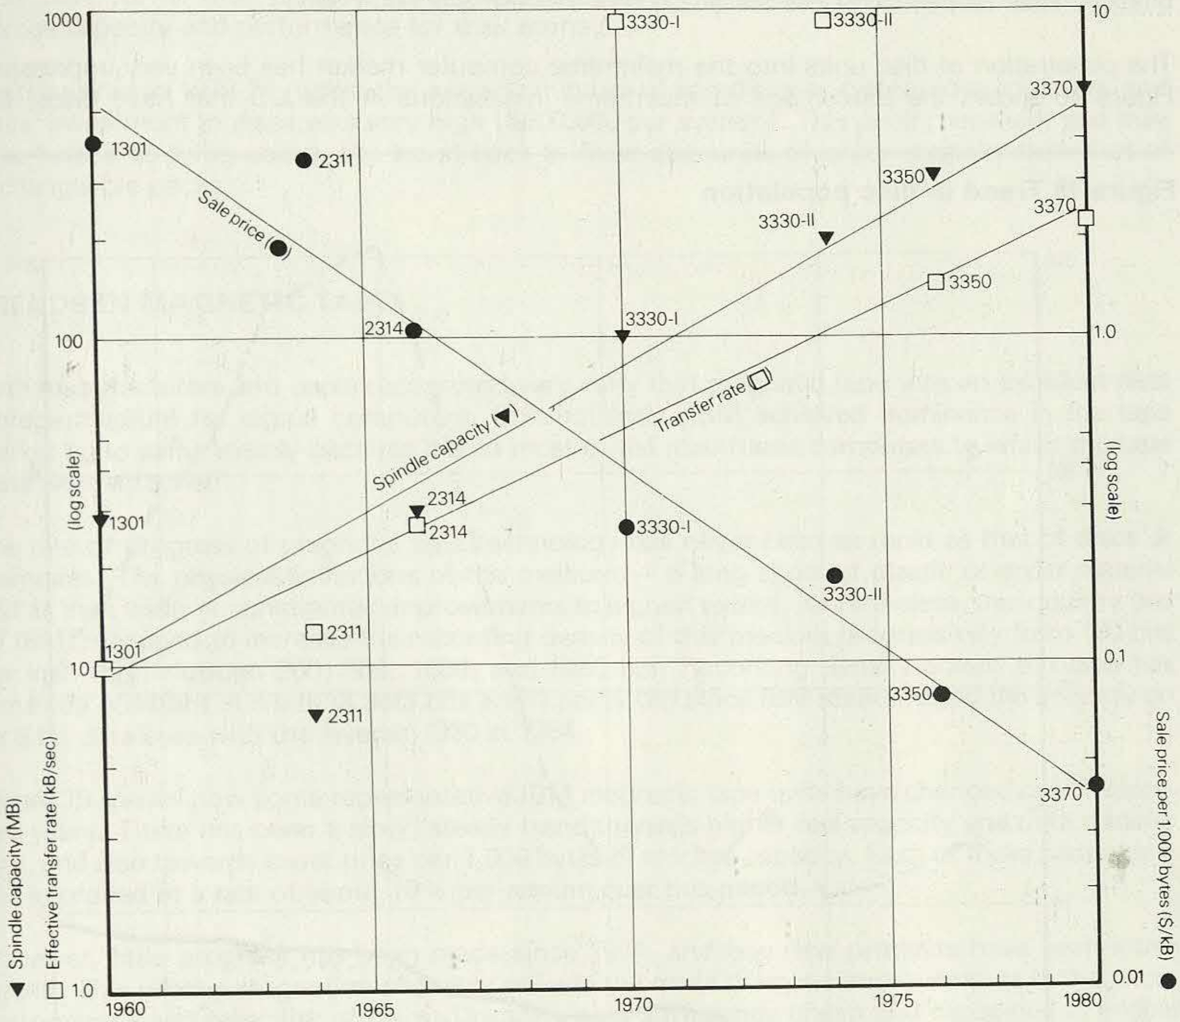

The three sets of points show the trends over twenty years in the price per thousand bytes, spindle capacity, and effective transfer rate of representative IBM disc units. The regression lines show the impressive trends in all three categories. Reliability has also greatly increased over this period.

(Source: Auerbach)

It is noteworthy that IBM has been the clear market leader in disc devices since the early 1960s. Although plug-compatible devices have competed with IBM products on price, no other manufacturer has been as innovative or as competitive as IBM in bringing out new disc products and in pricing them well within reach of the average mainframe user.

In the last few years, there has been considerable activity at the low end of the disc market as a result of the proliferation of inexpensive minicomputers. A scramble has developed amongst the independent disc suppliers to produce and sell an eight-inch multi-megabyte disc unit for just a few thousand dollars, a product that would have been unimaginable a decade ago.

At the high end of the mainframe disc market, IBM has introduced the 3370, with a capacity of more than 570 megabytes per spindle and a recording density of 7.5 million bits per square inch. This contrasts with the old 1301, which had a capacity of 28 MB and a storage density of 25,000 bits/in^2. Their respective prices are $62,000 (1979) and $42,000 (1961).

The penetration of disc units into the mainframe computer market has been very impressive. Figure 18 shows the percentage of mainframe installations in the US that have discs, the

Figure 18 Trend in disc population

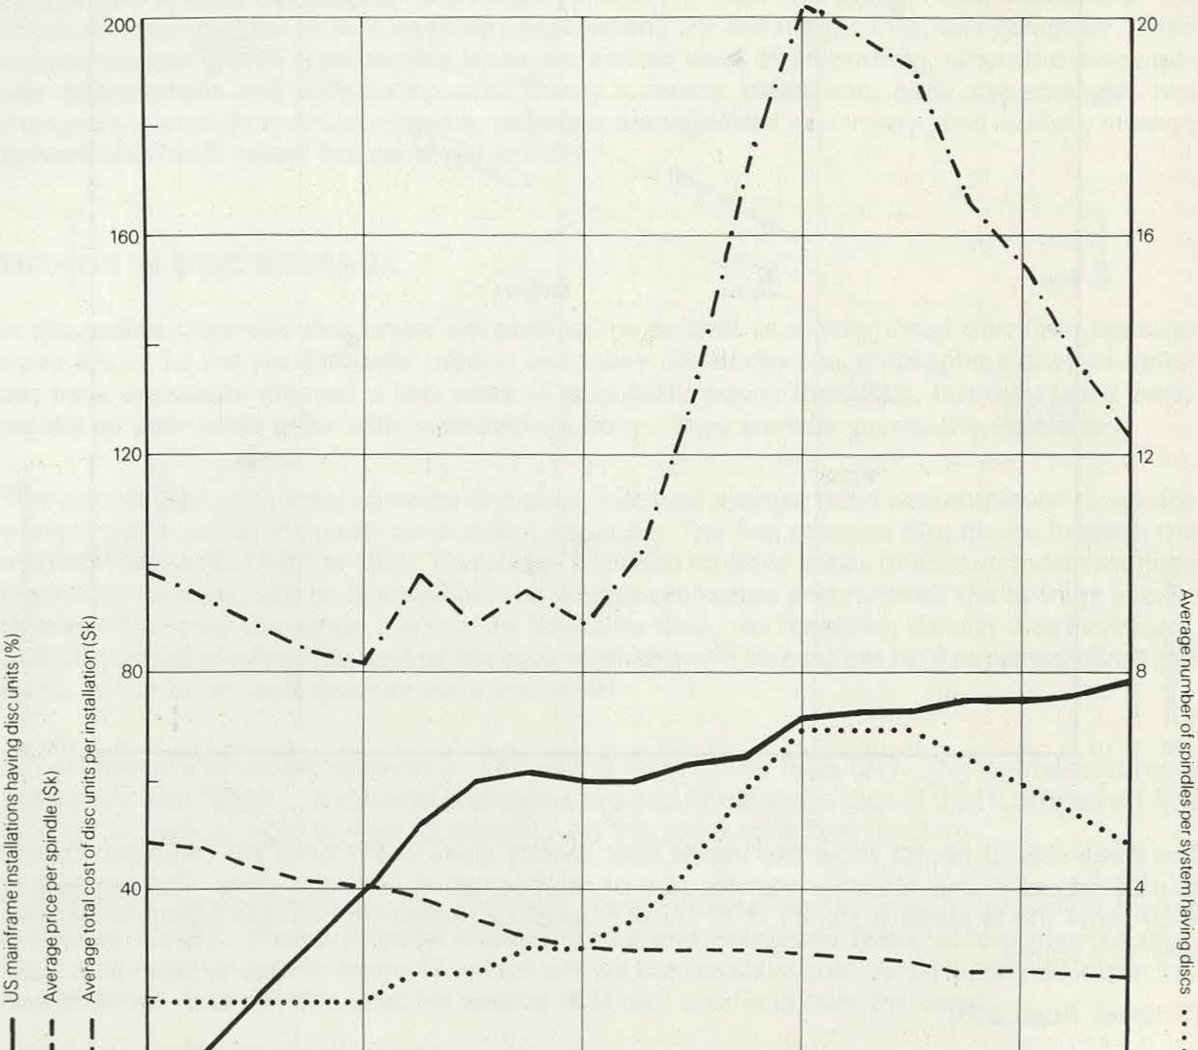

The curves show a high penetration of disc units in mainframe installations in the US. The number of spindles per system has declined somewhat since its 1972 maximum, reflecting the increase in capacity of single spindles and a trend away from exchangeable media. The average price per spindle has dropped consistently and dramatically, and total user expenditure on discs has declined with it since 1972.

(Sources: Phister and IDC)

average number of spindles per system, and their average cost per system to mainframe users. The average price per spindle (also shown) has declined sufficiently since 1960 (by half) to offset the increase in the number of spindles per system. Users therefore paid, on average, only about 24% more for all their discs in 1978 than they paid in 1960, and they received vastly more storage capacity and performance for their money.

Costs peaked in 1972-74, when the average number of spindles per system rose to seven, and users' investment in discs was very high ($200,000 per system). This peak preceded, and may have helped to bring about, the trend back to fixed-disc units of larger capacity than that of exchangeable packs.

TRENDS IN MAGNETIC TAPE

Both manufacturers and users recognised very early that magnetic tape was an excellent data storage medium for digital computers. Consequently, IBM achieved dominance in the tape market quite early, mainly because it sold most of the mainframe computers to which the tape units were attached.

The rate of progress of magnetic tape technology has never been as rapid as that of discs or memories. The physical limitations of the medium — a long spool of plastic or mylar material half an inch wide — constrained improvements to a great extent. Nevertheless, the industry (led by IBM) managed to increase the recording density of this medium progressively from 100 bits per inch (bpi) through 200, 800, 1600, and 6250 bpi. Recording density across the tape has remained constant at 9 bits (8 data bits and 1 parity bit) since IBM standardised the industry on an 8-bit data byte with the System/360 in 1964.

Figure 19 shows how some representative IBM magnetic tape units have changed over twenty-five years. There has been a slow, steady trend towards higher reel capacity and data transfer rate, and also towards lower price per 1,000 bytes of storage capacity. Each of these parameters has improved at a rate of about 10% per annum over this period.

However, little progress has been made since 1974, and few new products have been introduced. This relative stagnation probably reflects the more dramatic improvements in the price-performance and reliability of discs. Discs are now sufficiently cheap and capacious to enable new systems to be designed utilising *only* disc files, without magnetic tapes for either file backup or master file retention. Consequently, magnetic tapes are gradually entering a state of relative obscurity.

Figure 20 shows the pattern of mainframe user expenditure in the US on magnetic tape units over the period covered by this study. The tape's penetration of the mainframe market has increased as unit prices have declined. This growth was probably obtained at the expense of card-based data storage in very small mainframe systems. The number of magnetic tape units per system having tapes has remained remarkably constant over twenty years, varying between 4.5 and 4.9 tape units per system. In 1978 it stood at 4.6, and it had remained constant for six years before that.

THE UTILISATION OF DISCS AND TAPES

It is apparent from the data presented in this section that a marked shift in the use of discs and tapes has come about in the last ten years. Whereas, prior to 1970, most computer-based data tended to reside on magnetic tape, those systems that have been designed more recently have emphasised all-disc random-access files to the exclusion of tapes. Magnetic tapes, it seems, have been relegated to the following three roles:

1. As backup for disc master files, a role which is becoming decreasingly common.

2. As a medium for very large files, updated infrequently or having a high hit-rate of updated records when processed.

Figure 19 Trends in tape price, capacity and speed

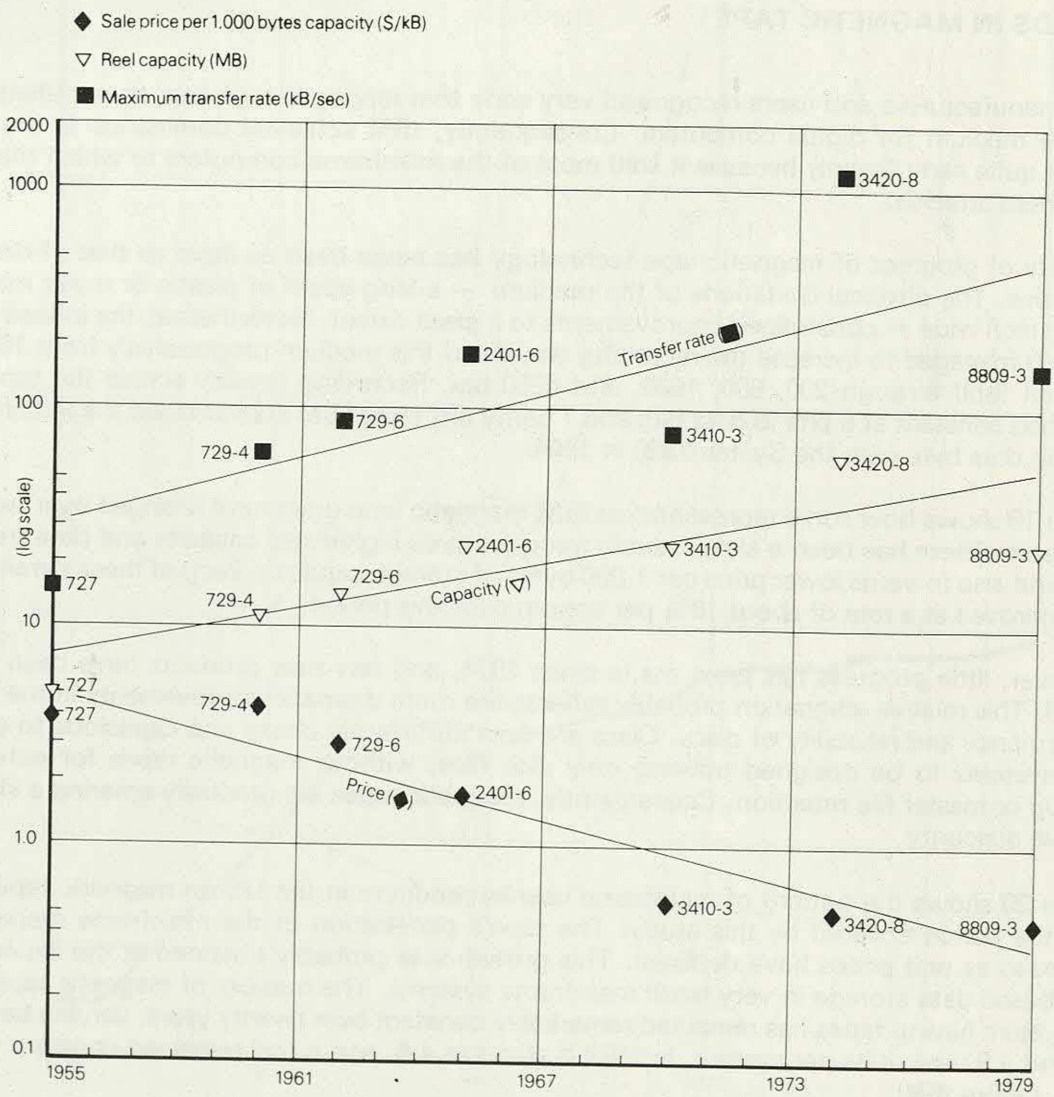

The points show the trends over twenty-four years in price per thousand bytes stored, spindle capacity, and maximum transfer rate of representative IBM magnetic tape units. The regression lines show the trend of each of these variables. Although the trends are well marked, they are extremely modest compared with those of disc drives.

(Sources: Phister and Auerbach)

3. As a vestige of old, tape-based systems, which would be too expensive to convert to disc during their remaining life.

Figure 21 provides some supporting evidence for this conclusion. The data displayed shows the total on-line and off-line data storage capacity for discs and tapes for all mainframe systems in the US over twenty years. In 1970, for the first time, the on-line disc capacity surpassed the on-line magnetic tape capacity, and by 1978 it was nine times larger. In fact, over the period 1968 to 1978, on-line disc capacity grew at an average rate of 46% per annum, while on-line tape capacity grew at a rate of only 10% per annum, which is less than the rate of growth of the mainframe computer market.

Figure 20 Trend of magnetic tape population

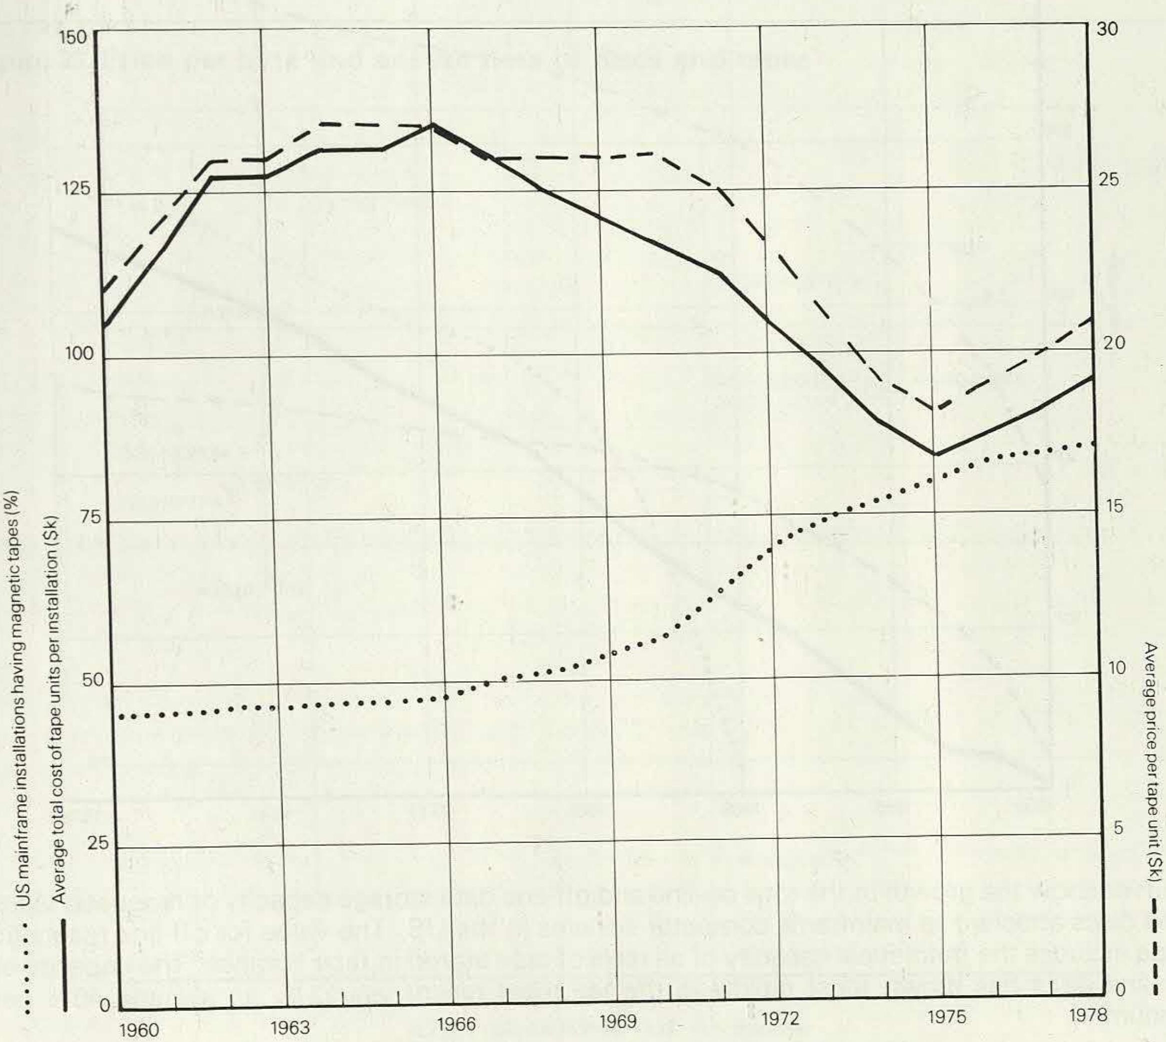

The curves show the increasing percentage of mainframe installations in the US having magnetic tapes, and the cost of those tapes to the user. The trend of use is high, but both unit price and total cost to the user have declined since 1970, an unusual combination.

(Sources: Phister and IDC)

Figure 21 Total storage capacity of discs and tapes

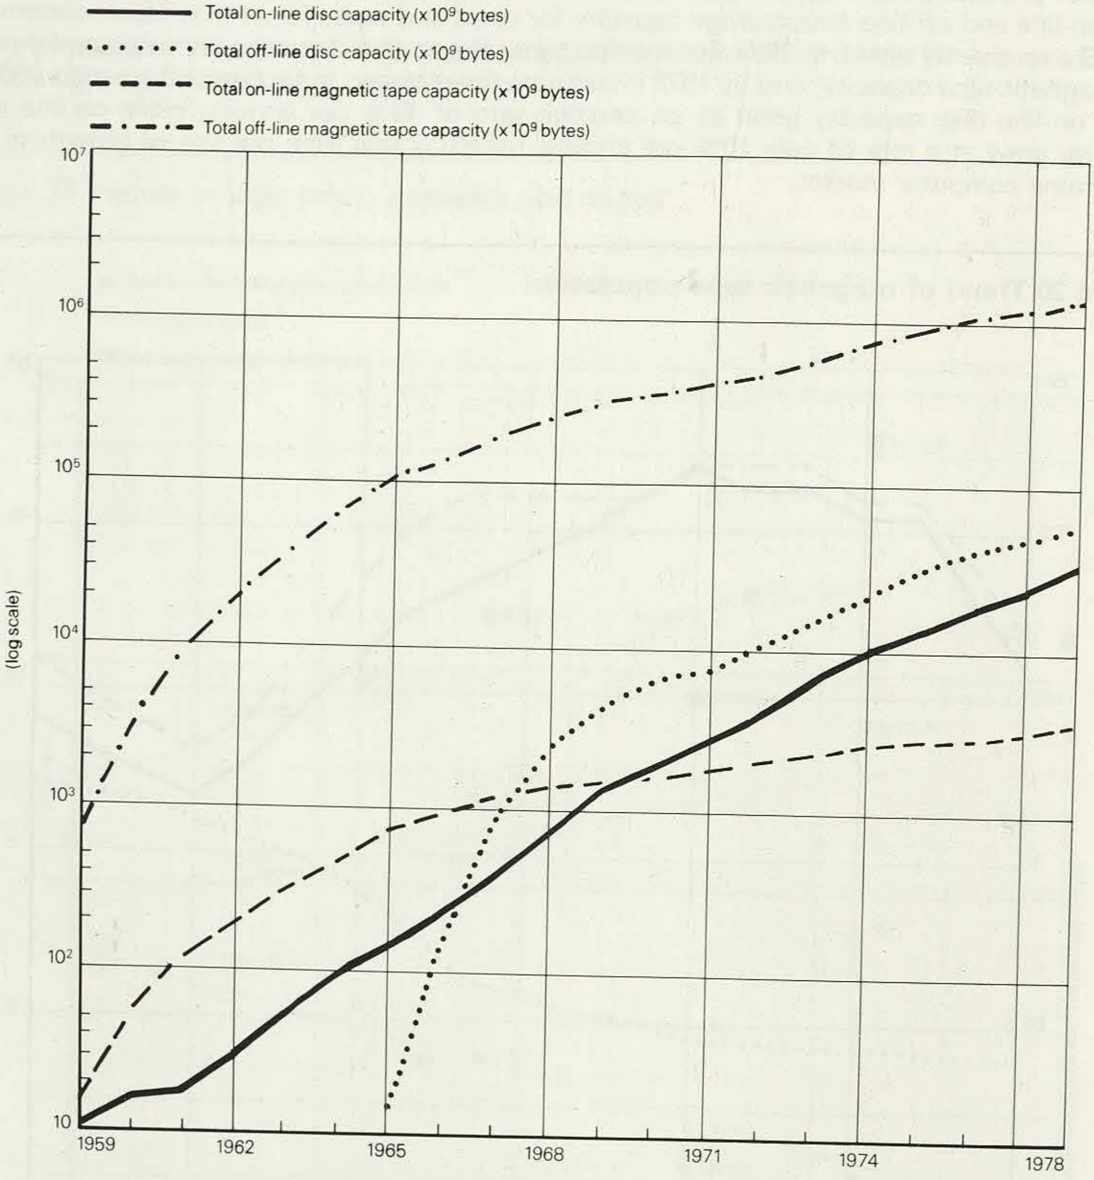

Curves show the growth in the total on-line and off-line data storage capacity of magnetic tapes and discs attached to mainframe computer systems in the US. The value for off-line magnetic tape includes the theoretical capacity of all reels of tape stored in tape libraries. The capacity of on-line discs has grown most rapidly in the ten most recent years, by an average 46% per annum.

(Source: Phister)

The total data storage capacity of all disc and tape media (both on-line and off-line) was about 1.55×10^{15} bytes in 1978, and was growing at 19% per annum. This stupendous capacity is difficult to comprehend. If that amount of data were recorded on magnetic tape at 1600 bpi, with no gaps or interruptions, the resulting tape would reach around the earth at the equator more than 600 times. If even a fraction of this capacity is actually occupied by users' data, we can only wonder whether the benefits that this information confers on users are commensurate with the quantity.

Finally, figure 22 shows the characteristics of discs and tapes compared with one another. The points represent the price per 1000 bytes ($/kB) and the access time (or rewind time) for several popular IBM discs and tape drives spanning twenty years. The slopes of the two regression lines show clearly the rapid development in both price and access time for discs, compared with that of tapes. By 1970, the absolute price per thousand bytes for disc storage was lower than that of tapes. There then followed the rapid increase in users' investment in discs, which was shown in figure 18. Since then, random-access storage devices have played a dominant role in system design and implementation, and tapes have diminished in importance to the point where they occupy only a supporting position today.

Figure 22 Price per byte and access time of discs and tapes

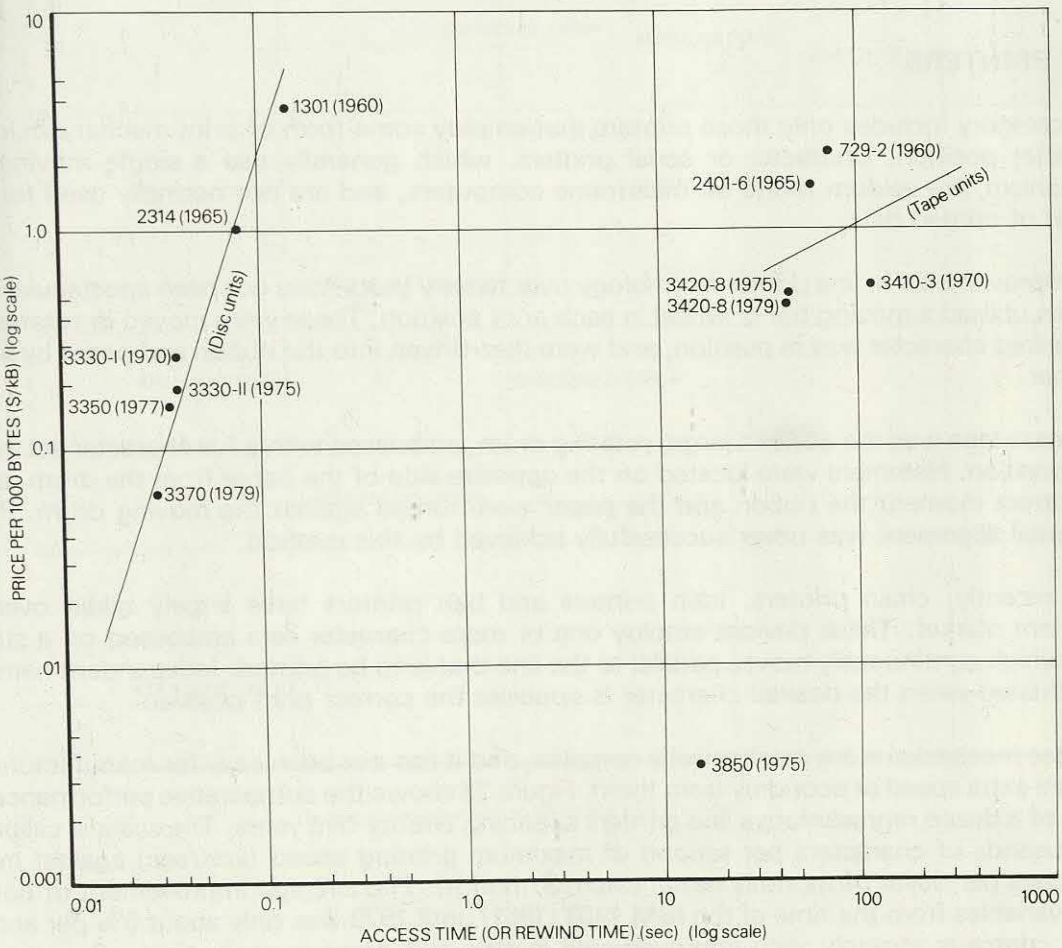

Points show the access time or rewind time (x scale) and the purchase price per 1000 bytes (y scale) for a selection of IBM disc and tape drives spanning the period 1960-79. Access time has improved little, but price per byte has declined dramatically in the case of discs, but less for magnetic tapes. The two regression lines show the separate trends very clearly.

(Source: Auerbach)

CHAPTER 4

OTHER PERIPHERALS

In addition to magnetic tapes and discs, numerous other peripheral devices are normally found in close proximity to mainframe computers. A subjective impression is that the variety of such devices is decreasing over time compared with, say, the mid-1960s. Users appear to be standardising on a few accepted types of peripherals.

This chapter examines four important classes of peripherals: line printers, unit-record equipment, data entry equipment, and terminals.

LINE PRINTERS

This category includes only those printers that employ some form of print mechanism in *each* character position. Character or serial printers, which generally use a single moving print mechanism, are seldom found on mainframe computers, and are not normally used for mass output of printed data.

The improvements in line printer technology over twenty years have not been spectacular. Early printers utilised a moving bar or wheel in each print position. These were moved or rotated until the desired character was in position, and were then driven into the ribbon and paper by a small hammer.

The next stage was the use of a large, rotating drum, embossed with a full character set in each print position. Hammers were located on the opposite side of the paper from the drum, and at the correct moment the ribbon and the paper were forced against the moving drum. Proper horizontal alignment was never successfully achieved by this method.

More recently, chain printers, train printers and belt printers have largely taken over this important market. These devices employ one or more character sets embossed on a strip or loop, which continuously moves parallel to the line that is to be printed. Independent hammers are actuated when the desired character is opposite the correct print position.

All these mechanisms are mechanically complex, and it has not been easy for manufacturers to squeeze extra speed or economy from them. Figure 23 shows the comparative performance and prices of a dozen representative line printers spanning twenty-five years. The scale is calibrated in thousands of characters per second of maximum printing speed (kch/sec) against million characters per dollar of monthly rental (Mch/$/month). The average improvement of both of these variables from the time of the IBM 1403 (1961) until 1979 was only about 5% per annum, which contrasts strongly with improvements in disc and memory technology. Progress has been steady, but modest by the standards of the computer industry.

However, there is a clear trend in the actual usage of line printers in mainframe installations in the US and in users' expenditure on them. Figure 24 charts the average printer price and the average number of line printers per mainframe system having printers. These show a clear trend downwards in the price of printers in use until recently, and a steady trend upwards in the number of printers per system from 1 to 1.8 in 1978.

Figure 23 Representative line printer prices and performance

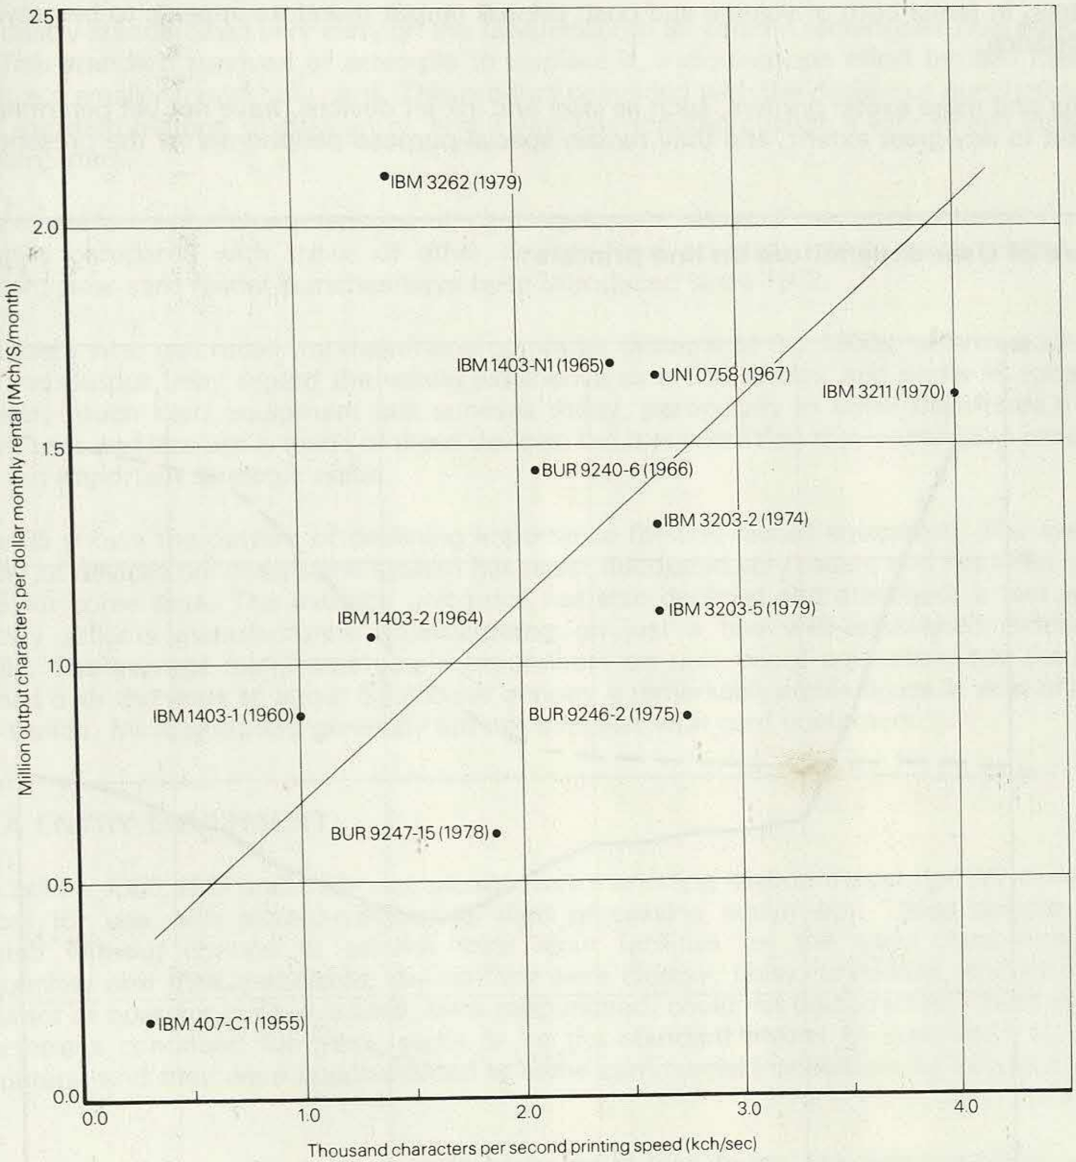

The points represent twelve popular line printers over twenty-four years. Each point shows the speed and price-performance of a printer. The regression line shows how these two variables are related, and which printers were the fastest and the best cost-performers. The further to the right and the higher up a point is, the better its performance.

(Source: Auerbach)

The remaining line on figure 24 shows the average annual expenditure by mainframe users in the US on line printers and printer controllers. This annual figure was obtained by first combining the purchase prices of the average number of printers per system and of one controller, dividing by 44 to obtain a notional monthly rental, and multiplying the result by 12 to produce an annual rental. This trend shows that user expenditure on printers started high, declined to a low of $13,200 per annum in 1970, and then doubled to $26,000 per annum by 1978, the increase being accounted for both by more printers per system (1.8 as against 1.2 in 1970) and by dearer printers. In terms both of volume and cost, printed output therefore appears to be growing in importance.

Newer and more exotic printers, such as laser and ink-jet devices, have not yet penetrated the market to any great extent, and they remain special-purpose peripherals for the present.

Figure 24 User expenditure on line printers

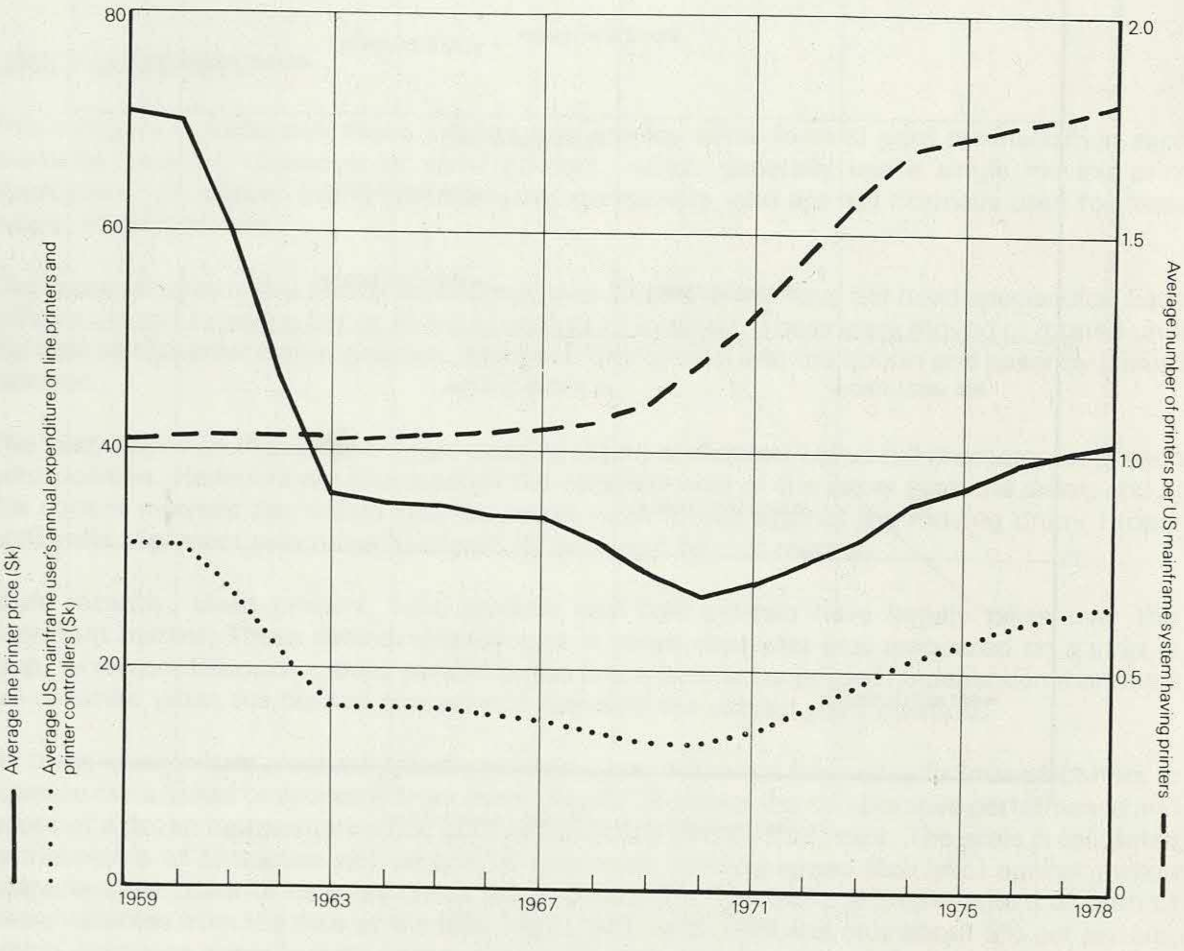

The curves track the movement of line printer prices in mainframe computer installations in the US and the average user expenditure on line printers in those installations. Also shown is the average number of line printers per installation, the growth of which has largely accounted for the increased overall expenditure.

(Source: Phister)

UNIT-RECORD EQUIPMENT

Unit-record equipment consists of on-line card-oriented devices, particularly card readers and punches. The performance of such equipment has been less affected by improving technology than has any other class of hardware. The punch mechanism has hardly changed at all in twenty years, whilst card readers underwent only one major transformation in that period, from wire brushes to photoelectric hole-sensing equipment.

The industry standardised very early on the IBM/Hollerith 80-column rectangular-hole punched card. This standard survived all attempts to displace it, including one effort by IBM itself to introduce a smaller round-hole card. This product coincided with the decline of punched cards as the standard method of data entry, and it was not adopted outside a very limited range of IBM computers.

The cost-performance characteristics of card equipment showed many inconsistencies and anomalies compared with those of other devices, and no clear trends were evident. No important new card reader-punches have been introduced since 1972.

Many users who can recall the mainframe computer systems of the 1960s, with their all-card input and output, may regard the whole experience as a bad dream, and prefer to forget it. However, much card equipment still survives today, particularly in small mainframe installations. Less and less use is made of these devices, but it appears that they continue to occupy a small but important strategic niche.

Figure 25 shows the pattern of declining importance for unit-record equipment. The average number of devices per mainframe system has never fluctuated very much, and has been stable at 0.8 for some time. The average unit price has also declined and stabilised, a fact which probably reflects manufacturers' standardising on just a few well-established production models. The average mainframe user's expenditure on unit-record equipment has therefore declined over the years to about $4,800 per annum, a remarkably stable figure in view of other price trends. Minicomputers generally are not supplied with card equipment.

DATA ENTRY EQUIPMENT

Even before 1950, IBM and other companies were marketing keyboard-operated punched card devices for use with electro-mechanical data processing equipment. These devices were adopted without change to provide data input facilities for the early computers. The keypunches and their associated key-verifiers were clumsy, noisy, unreliable, and extremely intolerant of operator error — a card, once mispunched, could not be corrected. These devices nevertheless continued for many years to be the standard means of data entry for most computers, and they were supplemented in some commercial installations by punched paper tape.

Mohawk Data Sciences then introduced a keyboard-to-tape device, the output of which was an IBM-compatible magnetic tape. Although the device was much more expensive than a keypunch, it offered vastly more flexibility and improved operator productivity. It was a commercial success, and it spawned a number of competitors.

In the late 1960s, the appearance of the cheap minicomputer made possible a multi-keyboard system, which accepted data through visual display units and stored it on a formatted disc. The first of these to achieve commercial success was the Computer Machinery Corporation's CMC-9, and, like the key-to-tape device, the new product's success generated much competition.

IBM replied rather feebly in 1970 with a buffered keypunch, but then struck back firmly in 1974

Figure 25 Unit-record equipment in mainframe systems in the US

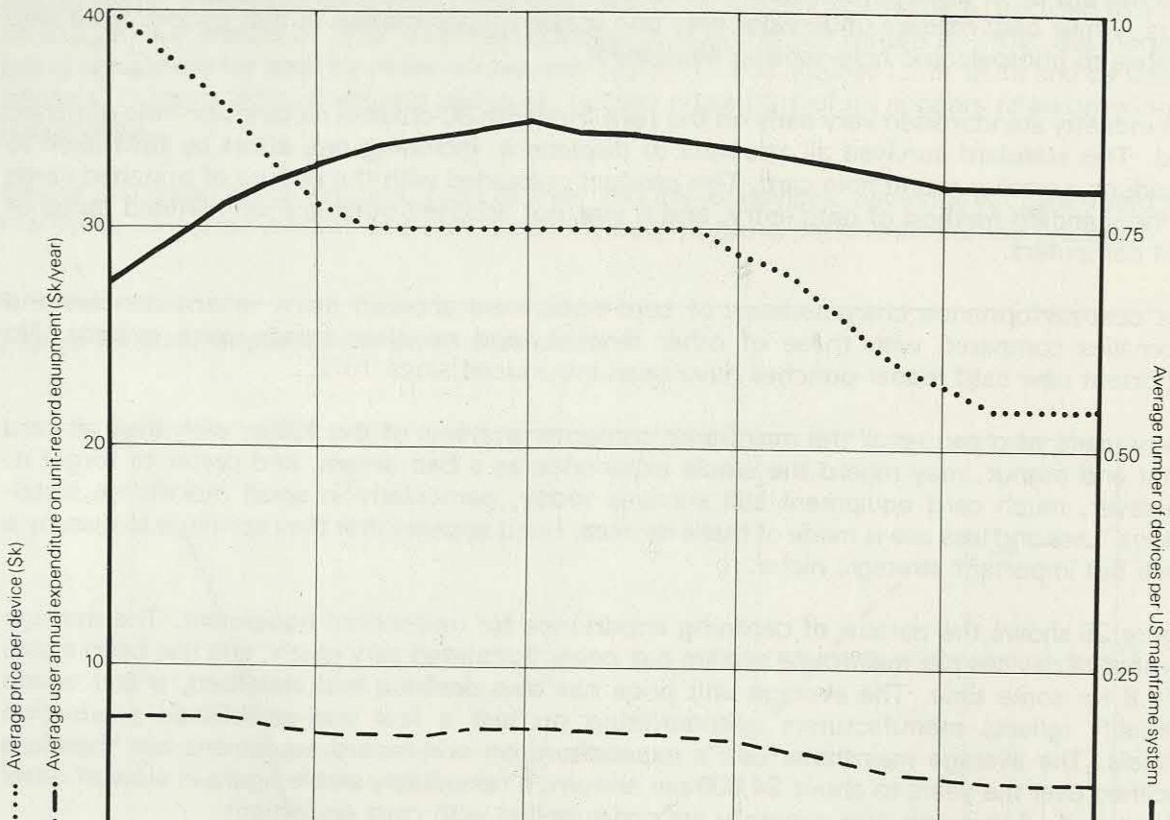

The curves show the declining use of unit-record equipment in mainframe computer installations over twenty years. The number of such devices per system has declined from a maximum of 0.87 in 1967 to a stable value of 0.80 by 1975 and beyond. Meanwhile, the unit price has declined considerably, reflecting the manufacturers' concentration on well-established but older models. Consequently, user expenditure on unit-record equipment has declined as well. This is one of the few categories of computer hardware in which an absolute decline in expenditure has been observed.

(Source: IDC)

with the floppy disc (or diskette) key-entry unit, which had a major impact on the market. By today's standards, however, it is slow and inflexible by comparison with the better mini-based systems, which can perform a good deal of error checking and cross verification on input data.

Figure 26 shows the distribution of various types of keyboard in the US over twelve years. The average number of keyboards per mainframe computer has remained remarkably constant at between 6 and 7.

Figure 26 Keyboard types in use in the US

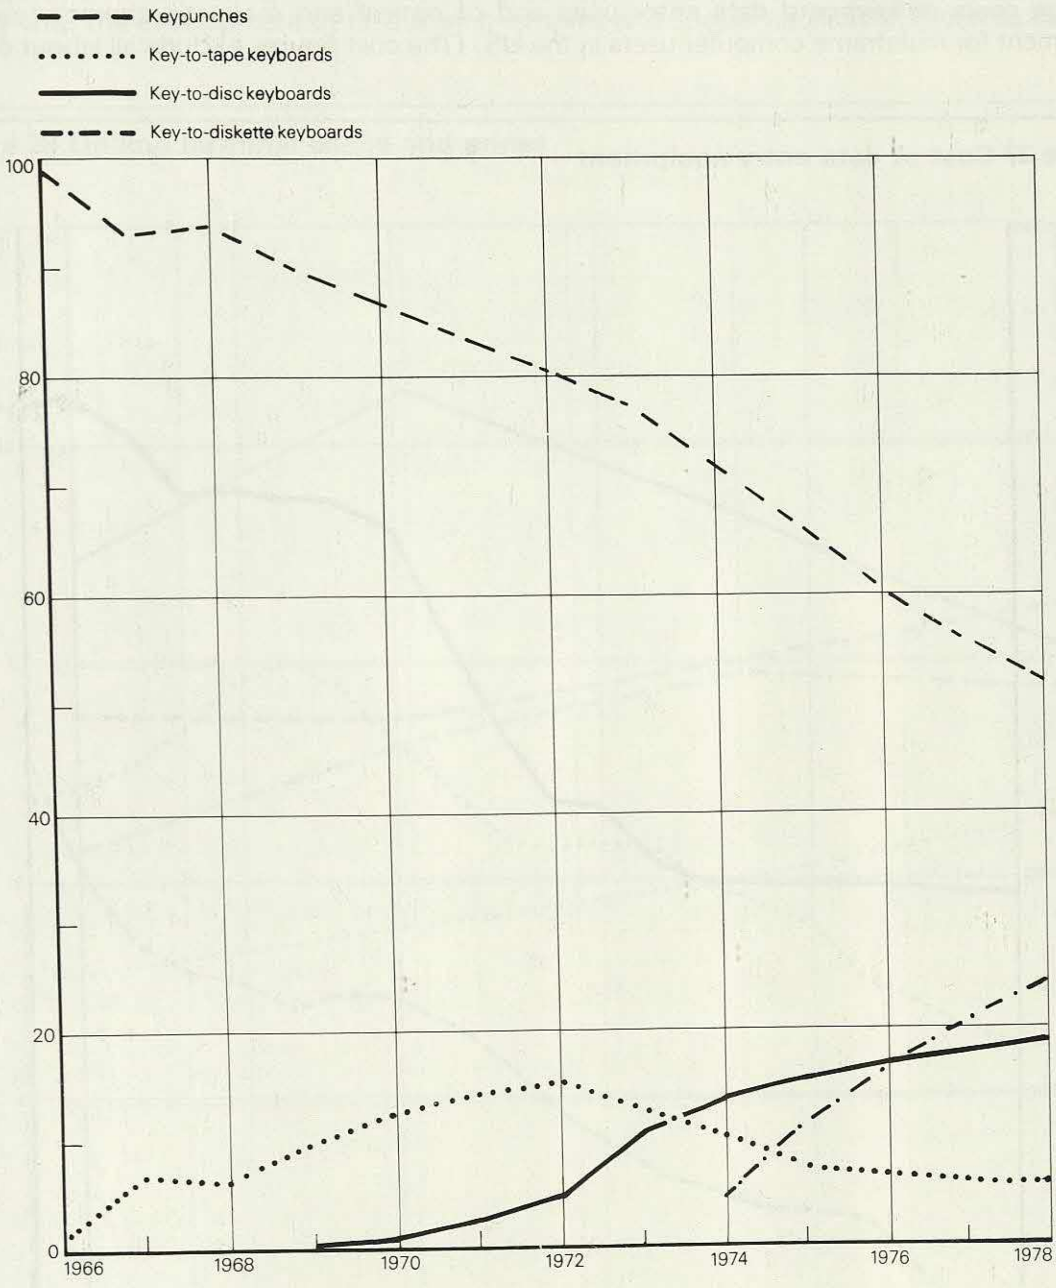

The figures show the changes since 1966 in the mix of keyboard types used in computer installations in the US. Each of the curves shows the percentage of the total keyboard population represented by one of the four categories. Keypunches, originally accounting for virtually all data entry keyboards, dropped to 52% by 1978, still a surprisingly high proportion considering the sophistication of more modern equipment. The key-to-diskette keyboards have shown the most rapid market penetration, capturing 24% of the total market in just five years since their commercial introduction by IBM.

(Source: IDC)

In summary, by comparison with electronics generally, this vital aspect of data processing has advanced at a very slow and uneven pace over the years. The result is that data entry costs have risen as computational costs have fallen. This trend is shown in figure 27, which displays the average costs of keyboard data entry units and of optical and magnetic character-reading equipment for mainframe computer users in the US. (The cost figures exclude all labour costs.)

Figure 27 Cost of data entry equipment

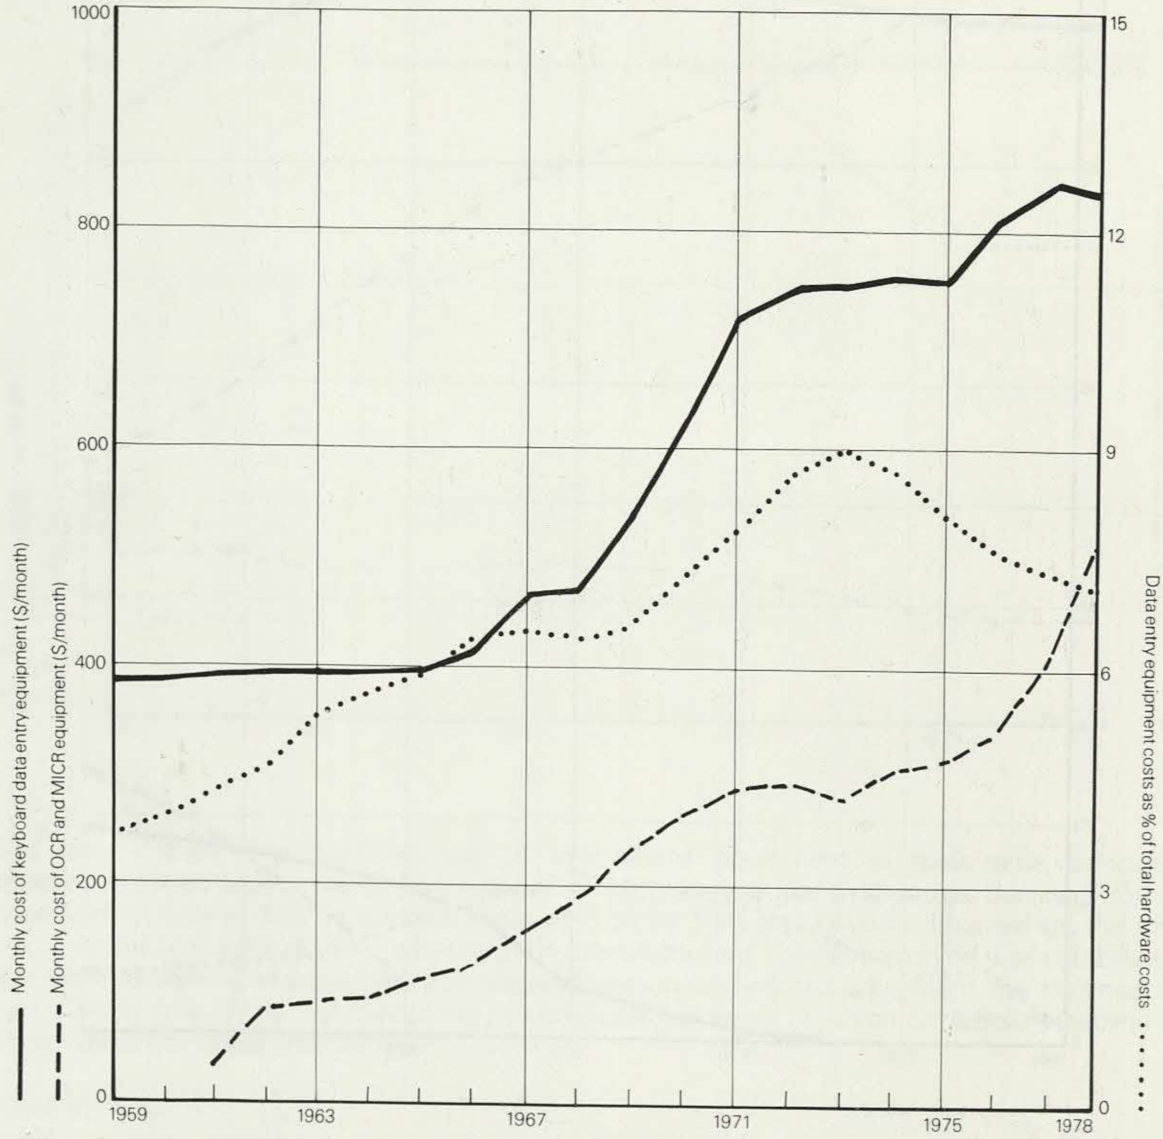

The figure shows the increasing expenditure on data entry equipment in mainframe computer installations in the US. Keyboard data entry systems still predominate, although some shift towards optical and magnetic ink character recognition systems has taken place. The cost of all this hardware has remained a fairly small fraction of total hardware costs, never exceeding 10% by value and declining to around 7% by 1978. This small fraction belies the importance of effective data entry equipment to most computer installations.

(Source: IDC)

TERMINALS

The use of on-line terminals with mainframe computers started in the early 1960s, and grew rapidly with the development of timesharing systems and special applications such as banking

Figure 28 On-line terminal usage and prices

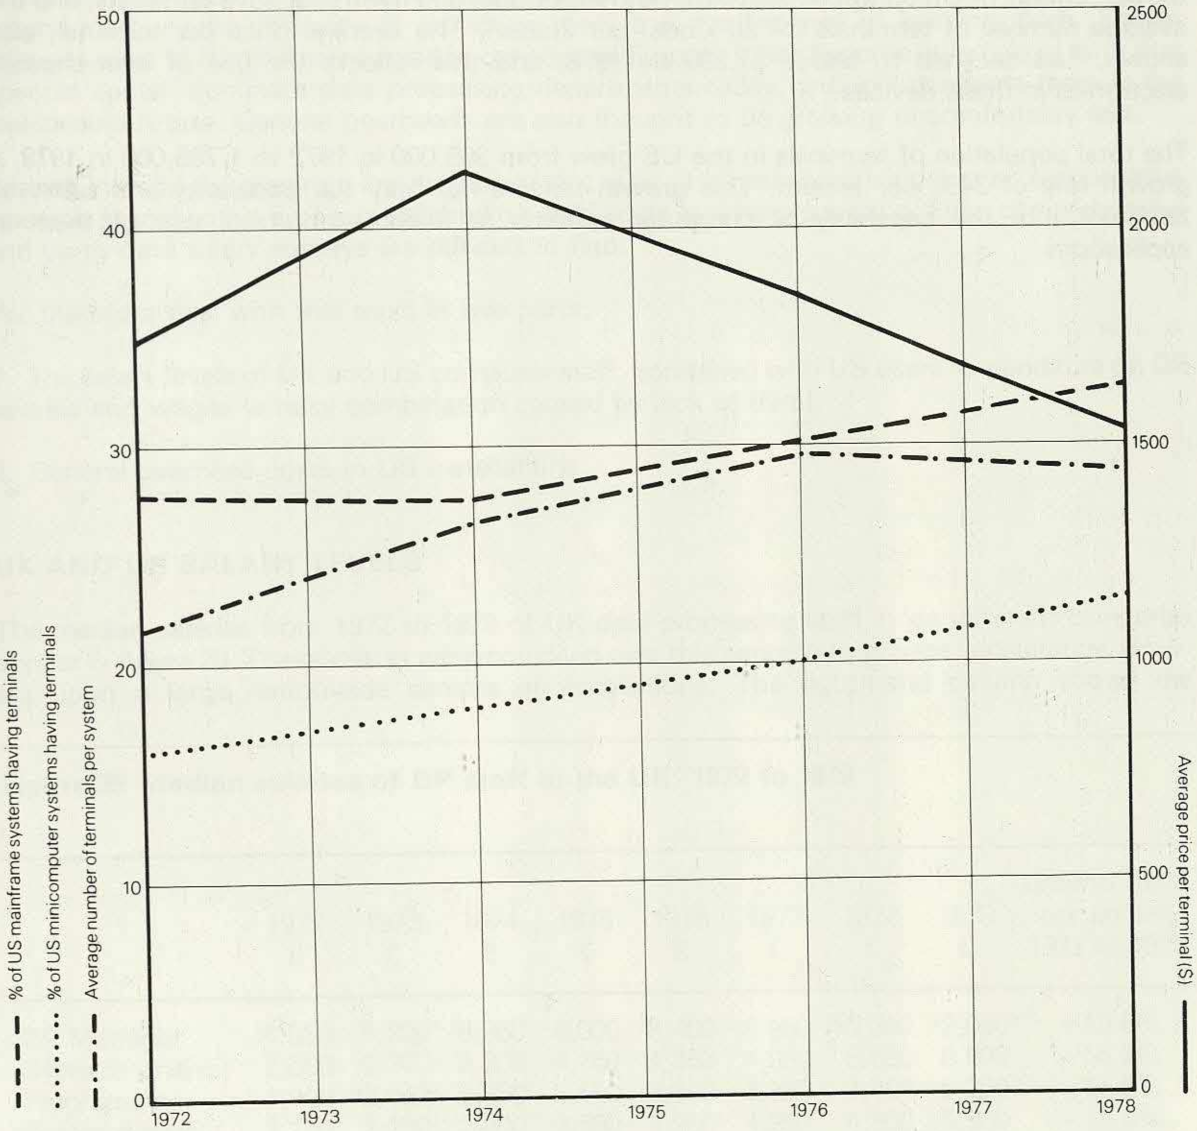

Two of these curves show the penetration of on-line terminals in the mainframe and minicomputer markets in the US since 1972. Despite the high average number of terminals per computer system having them (between 21 and 30), their percentage penetration was still not greater than a third of the total number of systems. Average prices of terminals have declined somewhat over this period, reflecting cheaper electronic components and higher production.

(Source: Phister)

and reservation systems. But the main impact of terminals was not felt until the widespread use of minicomputers introduced them practically everywhere in the computer world. Frequently, mainframe computers were left without terminals because of the difficulties their operating software encountered in handling on-line devices efficiently.

Like minicomputers, terminals are in the midst of a boom which has not yet showed signs of slowing down. The population of terminals in the US is roughly equally divided between printing terminals and visual display units, with a minority of special-purpose but higher-value terminals, such as those used in banking, point-of-sale, and share quotations. Figure 28 shows the percentage of US computer systems, both mainframe and mini, that have terminals, and the average number of terminals (of all kinds) per system. The average price per terminal, also shown, had declined to about $1,500 by 1978, and this reflects the use of ever-cheaper electronics in these devices.

The total population of terminals in the US grew from 305,000 in 1972 to 1,785,000 in 1978, a growth rate of 34% per annum. This growth reflects not only the popularity and utility of terminals, but the capability of computer systems to make economical use of them in applications.

CHAPTER 5

SALARIES AND GENERAL OVERHEADS

The greatest number of complaints voiced by data processing management in the UK centres around staff — the difficulty of attracting them and the expense of keeping them. Similar attitudes exist in North America and in continental Europe. It has become an article of faith that "people costs" dominate data processing departments today, and will dominate them in the foreseeable future. General overheads are also thought to be growing uncomfortably fast.

Bearing in mind the apparent importance of this topic, it is most surprising that no authoritative research seems to have been published on staff costs in data processing. Even comprehensive and up-to-date salary surveys are difficult to find.

We therefore deal with this topic in two parts:

1. The salary levels of UK and US computer staff, combined with US users' expenditure on DP salaries and wages (a risky combination caused by lack of data).

2. General overhead costs in US installations.

UK AND US SALARY LEVELS

The median salaries from 1972 to 1979 of UK data processing staff in six different categories appear in figure 29. These values were compiled over this period by a private consultancy, drawing upon a large nationwide sample of installations. The right-hand column shows the

Figure 29 Median salaries of DP staff in the UK, 1972 to 1979

	1972 £	1973 £	1974 £	1975 £	1976 £	1977 £	1978 £	1979 £	Growth trend per annum 1972 to 1979
DP Manager	4,950	5,200*	6,350	8,000	8,400	8,950	10,300	12,000*	+13.5%
Systems analyst	2,500	2,700	3,300	4,150	4,550	4,950	5,550	6,600	+14.9%
Programmer	1,900	2,050	2,550	3,150	3,500	3,850	4,400	5,200	+15.5%
Senior operator	2,150	2,250	2,800	3,650	3,950	4,250	4,950	5,950	+15.6%
Operator	1,700	1,850	2,250	2,900	3,250	3,600	4,150	4,950	+16.5%
Data prep operator	900	1,000	1,300	1,750	2,100	2,300	2,450	2,800	+17.6%
Average UK non-agricultural worker	1,863	2,128	2,529	3,098	3,482	3,790	4,389	5,138	+15.6%

* Estimated

(Sources: A private DP salary survey and the ILO Labour Statistics Yearbook) compound annual growth rate of each salary category over the whole period. The bottom line of data represents the average salary of all UK non-agricultural workers over the same period, obtained from a completely different source (the International Labour Organisation).

This data yields some very interesting observations:

1. The growth in DP salaries closely brackets the growth in wages generally and, in fact, the mid-point of all DP salary increases is exactly equal to the increase in general wages. This coincidence would have aroused our suspicions, had the data not come from completely separate sources.

2. A good ready reckoner seems to be that the salary of the job category "senior operator" is 10% to 15% above the UK average wage (particularly since 1976), and it is growing at about the same rate.

3. The lowest-paid staff showed the highest overall percentage increases. The best-paid staff fared worst, to the extent that programmers and above have not kept up with general wage increases in the UK.

4. Another way of viewing these salary relativities is that, compared with staff in general, the most expensive DP staff have actually become slightly cheaper to employ over this period.

5. The maximum salary differential — the ratio of a DP manager's salary to a data prep operator's — has been compressed from 5.5 in 1972 to 4.3 in 1979.

One may conclude from this data that the much-lamented expense of paying individual DP staff, at least in terms of their salaries relative to other workers, is illusory. The data suggests that they have barely kept pace with wage inflation at the lower levels, and in senior ranks they have actually fallen behind other workers.

However, this is only half the story, because it ignores the total costs of staff salaries at DP installations. It may be that average staff numbers have grown over time, causing the total wage bill (rather than individual wages) to grow faster than average wages.

SALARY COSTS AND STAFFING LEVELS

To begin this analysis, figure 30 presents a history of the average DP salary costs of mainframe computer installations in the US broken down by four categories of DP workers. (Americans tend to view the DP manager's salary as an overhead.)

What is immediately arresting is that the total wages of keyboard operators have, throughout the twenty years, been the highest of the four categories. Although the share of keyboard operators' wages of the total has fallen (from 36% to 28%) over the period, it is still the largest share. Furthermore, one suspects that this amount would be higher still but for the recent trend to employ non-DP staff in user departments to perform data input. As a rule, the wages of these workers are not accounted for in DP budgets. Hence the figures shown for keyboard operators may be artificially low in terms of the true costs to the company.

It is interesting too that computer operators' total wages are, by a substantial margin, the fastest growing of the four categories. This fact is difficult to reconcile with what we already know about changes in hardware. It leaves three questions that cannot be answered from the evidence available:

1. Have computer operators' individual wages increased faster than other wages? (British result: not much.)

2. Are computers (together with their operating systems) more difficult to operate in 1978 than they were in the 1960s?

3. Does the average mainframe installation have more computers (or more equipment requiring operator intervention) than it did in the 1960s?

Lastly, we note another very good ready reckoner from the American data: the combined total salaries of systems analysts and programmers are very nearly half the total of all DP wages. This relationship seems to hold true over the entire twenty years shown in figure 30.

Figure 30 Annual staff salary costs in mainframe installations in the US ($k)

Year	Systems analysts	Programmers	Computer operators	Keyboard operators	Total
1959	15.4	16.5	8.8	23.0	63.7
60	16.3	17.3	9.5	24.2	67.3
61	17.4	18.0	10.2	25.3	70.9
62	18.7	19.7	11.1	26.7	76.2
63	18.3	18.6	10.5	26.4	73.8
1964	17.9	17.9	10.2	27.4	73.4
65	19.0	19.6	11.0	29.7	79.3
66	21.5	22.5	12.2	32.0	88.2
67	26.8	27.3	14.8	37.6	106.5
68	30.1	31.1	16.5	39.0	116.7
1969	36.9	38.4	20.3	43.6	139.2
70	42.8	44.4	24.7	50.3	162.2
71	43.7	46.6	27.8	54.8	172.9
72	43.4	46.7	29.1	60.2	179.4
73	41.7	43.2	29.0	62.7	176.6
1974	40.8	44.7	30.3	67.6	183.4
75	46.1	51.3	34.9	73.7	206.0
76	57.6	64.3	43.2	83.5	248.6
77	71.8	78.1	52.4	91.2	293.5
78	86.2	93.4	62.6	96.2	338.4
20-year growth rate (1959 to 1978)	9.5% p.a.	9.6% p.a.	10.9% p.a.	7.8% p.a.	9.2% p.a.
10-year growth rate (1969 to 1978)	9.9% p.a.	10.4% p.a.	13.3% p.a.	9.2% p.a.	10.4% p.a.

(Source: IDC Surveys)

We next present some additional data on individual salaries compiled by *Datamation* magazine in its annual surveys. This data is contained in figure 31. We show only the most recent years of this survey (1973 to 1978), and the figures give the average individual annual salaries of DP staff in the categories shown.

Two immediate observations arise from this data:

1. The salaries are lower than many British estimates of the earnings of DP staff in the US.

2. The annual rates of increase, ranging from 5.4% to 6.3% per annum, are low for what in most countries was (and still is) a period of high wage inflation. In fact, the "official" increase in non-agricultural wages in the US over the same period was 6.9% per annum.

At the bottom of figure 31 this data is combined with the total wage expenses of figure 30 to determine the average number of staff per installation and the growth rates of staff numbers by category of worker. This analysis shows a total growth in staff numbers of 4% per annum over the period 1973 to 1977. (We have left 1978 out of this analysis because some of the salary values seem strangely low.) Leaving data prep operators aside, for reasons discussed above, then the overall increase in DP staff numbers is about 5% per annum over the period.

Figure 31 Average DP staff salaries and numbers in the US

Year	Systems analysts	Programmers	Computer operators	Keyboard operators
Salaries (in $k p.a.)				
1973	14.1	12.1	8.9	7.1
1974	15.3	12.8	9.4	7.4
1975	17.0	14.1	9.9	8.2
1976	17.8	14.7	10.5	8.3
1977	18.0	15.0	11.0	8.9
1978	18.4	13.8	10.7	8.9
Annual growth rate 1973 to 1977	6.3% p.a.	5.5% p.a.	5.4% p.a.	5.8% p.a.
Average number of staff per mainframe installation in the US				
1972	3.3	4.1	3.5	9.0
1977	4.0	5.2	4.8	10.2
Rate of increase	4% p.a.	5% p.a.	6.5% p.a.	2.5% p.a.

(Sources: *Datamation* and Figure 30)

We emphasise that combining the data from separate surveys in this way gives rise to low credibility in the results. We doubt whether the basic data of either figure 30 or figure 31 is particularly reliable in the first place, and so a combination of the two is a flimsy measure indeed. Nevertheless, if the basic data is consistent year by year, then the calculated rates of increase should be fairly consistent as well.

Finally, figure 32 shows the composition of staff of the average mainframe installation in the US in 1977, both by the number of staff of each kind and by the value of their total salaries. Our interpretation of these costs is given in the chapter on total system costs beginning on page 62.

Figure 32 Composition of DP staff by function, 1977

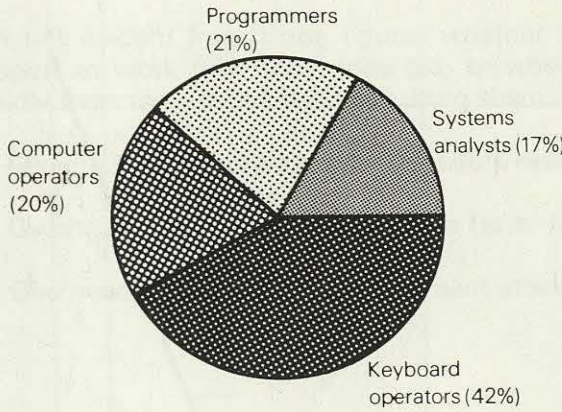

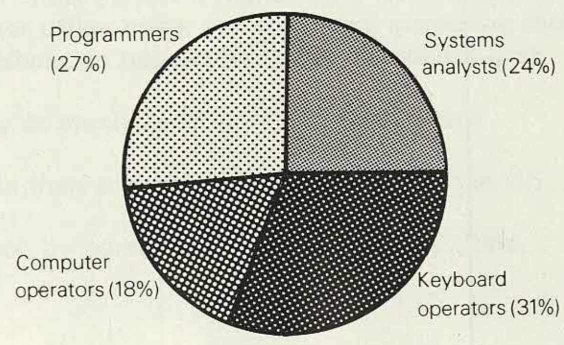

The two diagrams in the figure represent the composition of DP staff in the average mainframe computer installation in the US in 1977. The diagram on the left shows the percentage composition by number of staff; that on the right, by wages per category.

(Sources: *Datamation* and IDC surveys)

GENERAL OVERHEADS

Accounting practices differ somewhat between the various countries of Europe and the US in the way they measure and allocate general overheads in DP organisations. In order to make the best use of the available data, we have adopted the American convention of expressing overheads as a percentage of direct salaries and wages. The assumption behind this practice is that most overheads increase roughly in proportion to staff costs — an assumption which it is easy to dispute, but is probably a good rule of thumb for much of the time.

The overheads included in this category are the usual staff-related costs: pensions, fringe benefits, office space (for staff, not for computers), secretaries, telephones, office equipment, general supplies and — notably — DP management. All these costs and a few others are included in the heading of "overheads", and we do not attempt either to break them down further or to analyse the component costs separately.

Figure 33 shows the historical overhead costs of the average mainframe installation in the US both in money terms ($/year) and as a percentage of total staff salaries.

These results may surprise some managers. The average total overhead costs themselves rose

Figure 33 Trends in general overhead costs

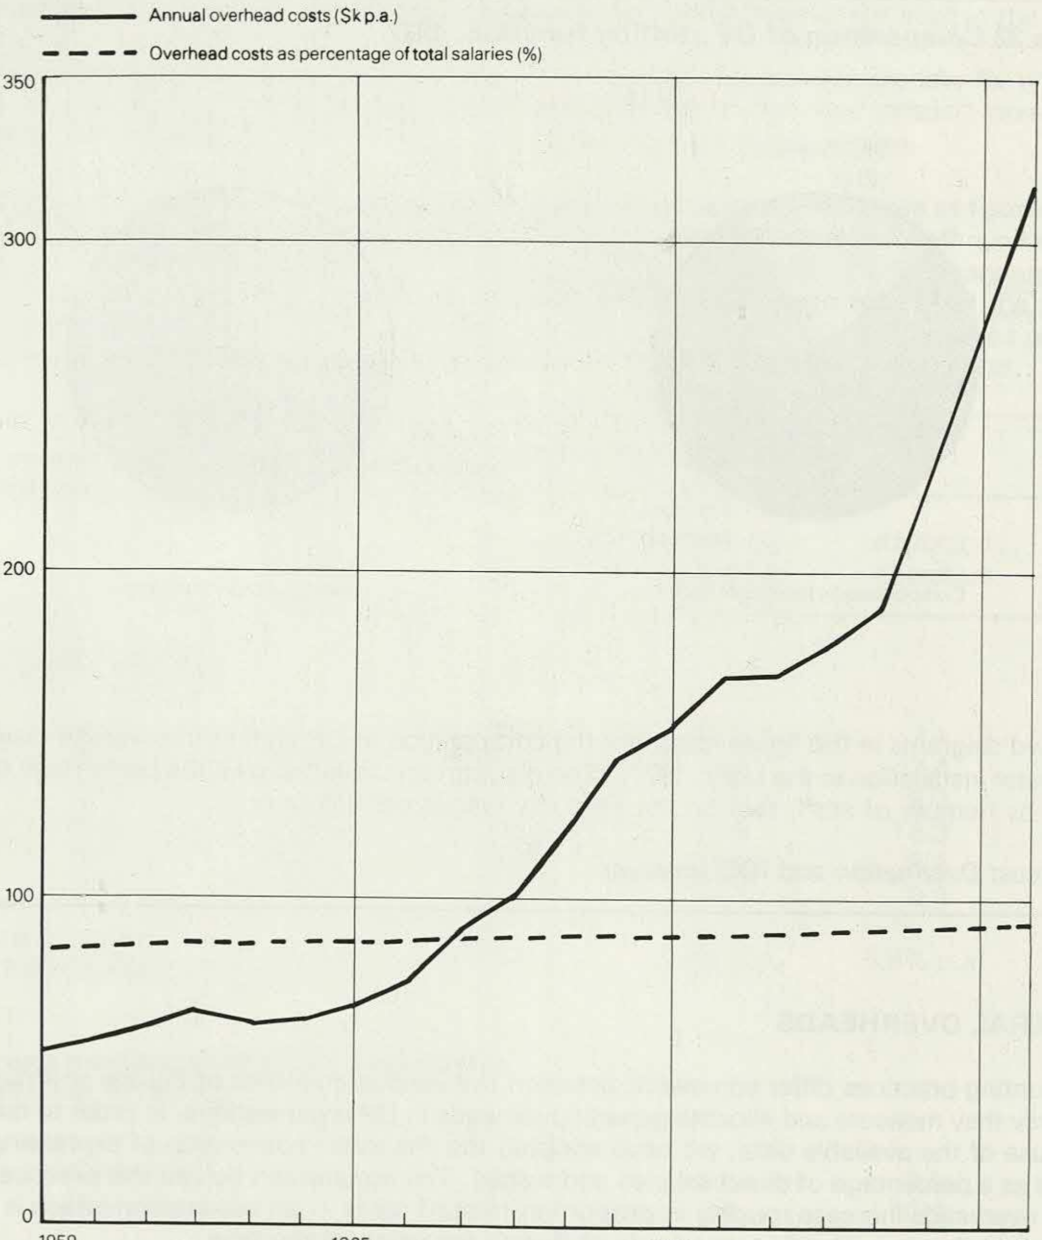

The curves show the trend in general overhead costs in mainframe computer installations in the US. These have shown a compound annual increase of nearly 10% p.a. over the twenty years covered by the figure, and an increase of 17.2% p.a. over the most recent period, 1974-1978. Overheads have also increased as a percentage of direct salaries, from 85% in 1959 to 94% in 1978.

(Source: Phister)

from $54,000 per annum in 1959 to $319,000 per annum in 1978, which was equivalent to an annual rate of growth of almost 10% per annum over the entire period. In recent years, the growth was much higher, and between 1974 and 1978, overhead costs grew by more than 17% per annum.

Even more surprising is that overheads expressed as a percentage of direct salaries (as inflation-proof a measure as could be hoped for) also increased, from 85% in 1959 to 94% in 1978. This means that, in addition to salaries, users spent almost as much again to support their staffs with administration, accommodation, and fringe benefits.

It is not evident from these figures whether DP staff expect a higher level of amenities and support at work than they once did, or whether office costs are somehow increasing more rapidly than the cost of the staff using them. What can be said from the evidence is that:

1. General overheads now cost DP users nearly as much as all staff salaries together.

2. Overhead costs are increasing at a faster rate than are staff salaries, at least in the US.

3. Overheads receive little management attention by comparison with staff salary costs.

CHAPTER 6

SERVICES AND FACILITIES

This chapter presents, in effect, an "all others" collection of DP costs that do not fit into any of the previous categories. This is not to say, however, that these are unimportant. The costs discussed are the following:

— Equipment maintenance

— Physical facilities

— Data processing supplies

— Purchased software and services

— Data communications services

From the above list, it is apparent that a particular mix of costs will never be the same for any two computer organisations. The average costs given should therefore be interpreted in this light, and they must be related to a non-existent "average" DP organisation, rather than to a particular user.

EQUIPMENT MAINTENANCE

Until very recently, computer users regarded maintenance charges either as an additional hardware cost or as an insurance policy against prolonged down-time. With the recent appearance of private maintenance companies, users have begun to regard maintenance charges as an additional source of revenue for computer manufacturers.

We have converted maintenance charges into a form in which they are comparable with one another across the entire spectrum of hardware. Their cost is expressed as the monthly charge per $100,000 purchase price of the equipment.

Figure 34 shows the trend in IBM's maintenance charges for each of five categories of equipment, and a weighted average for all units. Each of these lines tells an interesting story. As might be expected, the highest maintenance charges are imposed on those devices that have a high mechanical content, particularly line printers and card equipment.

However, it is processors and memory — the electronic devices — that have incurred the steepest percentage increases in maintenance charges since 1963, growing on average by 11% per annum. This increase should not be attributed entirely to inflation, because the measure we have chosen is reasonably inflation-proof. The fact is that IBM's maintenance prices on processors have increased much faster than inflationary factors warrant. This trend is particularly marked with the 370/1X5, 138, 148, 303X, and System/3 computers.

There are three possible explanations for this trend:

1. Processor reliability has deteriorated over the last fifteen years.

2. IBM's maintenance procedures and facilities are increasingly inefficient and costly.

3. IBM is increasing its profits on maintenance, perhaps so that it can appear to reduce hardware prices.

Figure 34 Trends in IBM maintenance charges

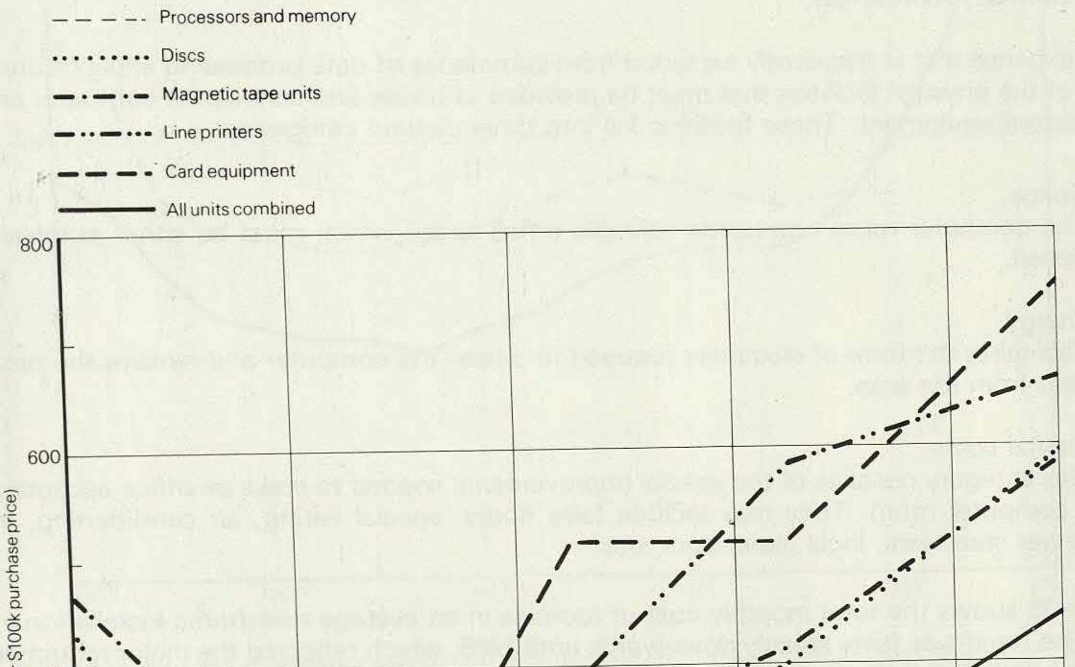

The six lines show the trend of IBM's maintenance charges for several classes of computer equipment. Values shown are the monthly maintenance charges (in $) per $100,000 purchase price of the equipment, for new units only. The most startling increases have been those for processors and memory, which are also presumably the most reliable and the easiest units to repair. Unit-record (card) equipment has taken over from magnetic tapes and line printers as the most expensive category overall.

(Source: Phister)

In the absence of detailed studies of this subject, and given the recent success of private maintenance companies, we are inclined to believe the third view.

We have not separated maintenance costs from the costs of hardware itself. Consequently, the equipment costs shown in the preceding chapters of this report include maintenance charges as well.

PHYSICAL FACILITIES

One expense that is frequently excluded from summaries of data processing expenditure is the cost of the physical facilities that must be provided to house and service the computer and the associated equipment. These facilities fall into three distinct categories:

1. Space
   The computer room represents valuable office space which must be either purchased or leased.

2. Energy
   This takes the form of electricity required to power the computer and remove the resulting heat from the area.

3. Capital costs
   This category consists of the special improvements needed to make an office acceptable as a computer room. They may include false floors, special wiring, air conditioning, mains power stabilisers, local generators, etc.

Figure 35 shows the total monthly cost of facilities in an average mainframe installation in the US. The trend was fairly steeply downwards until 1965, which reflected the major reductions in the size and the power consumption of computers brought about by second-generation (transistor) equipment. Since then, the trend was relatively stable until 1974, when installation size began to increase rapidly and facility costs followed.

The figure also displays the average floor space (in square feet) required by mainframe computer installations up to 1974, and the electrical power consumed by the installation (in megawatt-hours) per month.

The components of total facilities cost in American installations are very roughly as follows:

— Space rental	50%
— Electrical power	25%
— Capital costs (depreciation)	25%

DATA PROCESSING SUPPLIES

The cost of data processing supplies is a surprisingly high proportion of operating costs. The cost of paper-based media (printer forms and card stock) is particularly high. Because computers produce so much printout, there is a strong temptation for managers to consider all the cost to lie in producing the data, and to regard the paper itself as being free. This is far from true.

Figure 35 Facility costs of mainframe computer installations in the US

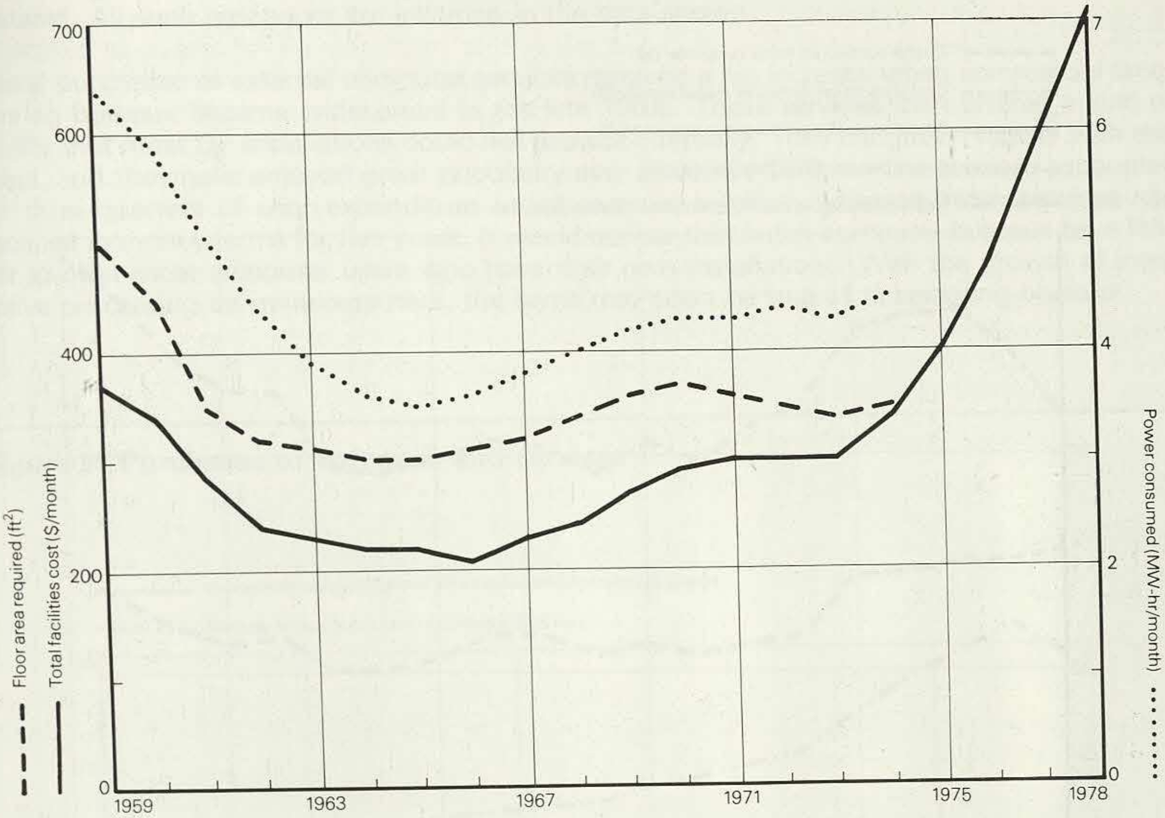

The curves show two components of facility costs, floor area and electrical power consumption, and the total monthly expenditure on facilities in mainframe computer installations in the US over twenty years. The interesting fact about all these variables is how little they have changed over a long period of time, considering the very great changes in the physical composition of the average computer system. Facility costs, like many others, show an unexpected sharp rise beginning about 1975.

(Source: IDC)

Figure 36 shows the average annual cost of supplies for each of the major peripheral devices: the cost of continuous forms per printer, of cards per output punch, of tape reels per magnetic tape unit, and of disc packs per spindle (all spindles counted). It also shows the total annual cost of supplies for the average mainframe computer installation in the US.

It is immediately apparent that the cost of printer forms is comparable to the annual cost of a printer itself. Furthermore, although the cost of tape reels and disc packs *per unit* is fairly low, their total cost to the installation must be fairly high, because most mainframe computers have multiple units of both types.

The total cost of DP supplies remained fairly steady at a low level between 1966 and 1973 but, like many other DP costs, it is now climbing rather steeply. In our experience, many users do not make a serious effort to control these costs, and most users are unaware of how much their installations spend on paper forms. This extravagance in printing often has its roots in poor system design.

Figure 36 Cost of computer supplies by type of device

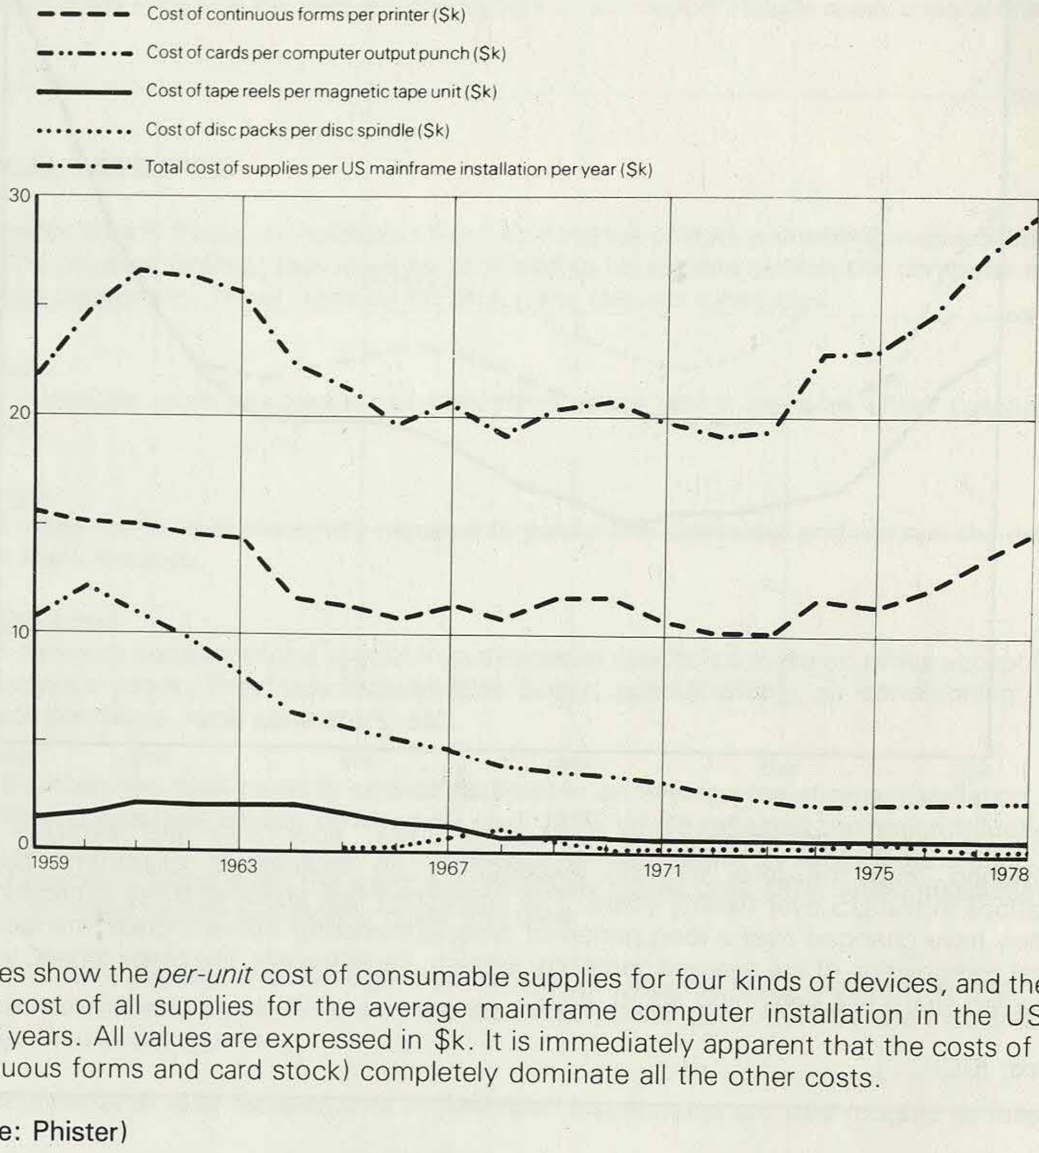

The lines show the *per-unit* cost of consumable supplies for four kinds of devices, and the total *annual* cost of all supplies for the average mainframe computer installation in the US over twenty years. All values are expressed in $k. It is immediately apparent that the costs of paper (continuous forms and card stock) completely dominate all the other costs.

(Source: Phister)

PURCHASED SOFTWARE AND SERVICES

This category of costs comprises expenditure on externally-prepared software and externally-provided technical services. In the former category fall mainly applications packages, system aids, and operating system software. In the latter, we include bureau services (both batch and on-line), contract staff, and consultancy.

Figure 37 shows the rapidly growing trend in expenditure on software purchases, and a slower trend in services. No information was available on the breakdown of software purchases, but from the pattern shown it would be safe to conclude that a good deal of this expenditure has gone towards "unbundled" IBM operating software. These programs were formerly offered to users without additional charge, but in the 1970s this practice ceased. There is also a growing willingness amongst users to purchase proprietary system aids, such as sort programs, teleprocessing monitors, system utilities, and (in larger installations) database management systems. All such purchases are included in the data shown.

Users' purchases of external computer services received a big increase when commercial timesharing bureaux became widespread in the late 1960s. These services then offered a kind of facility that most DP installations could not provide internally. They caught on rapidly with end users, and they have enjoyed great popularity ever since. By 1978, on-line services accounted for three-quarters of user expenditure on all external services, whereas batch services had declined in money terms for five years. It would appear that batch computer bureaux have little left to offer most computer users who have their own installations. With the growth of interactive processing on minicomputers, the same may soon be true of timesharing bureaux.

Figure 37 Purchases of software and services

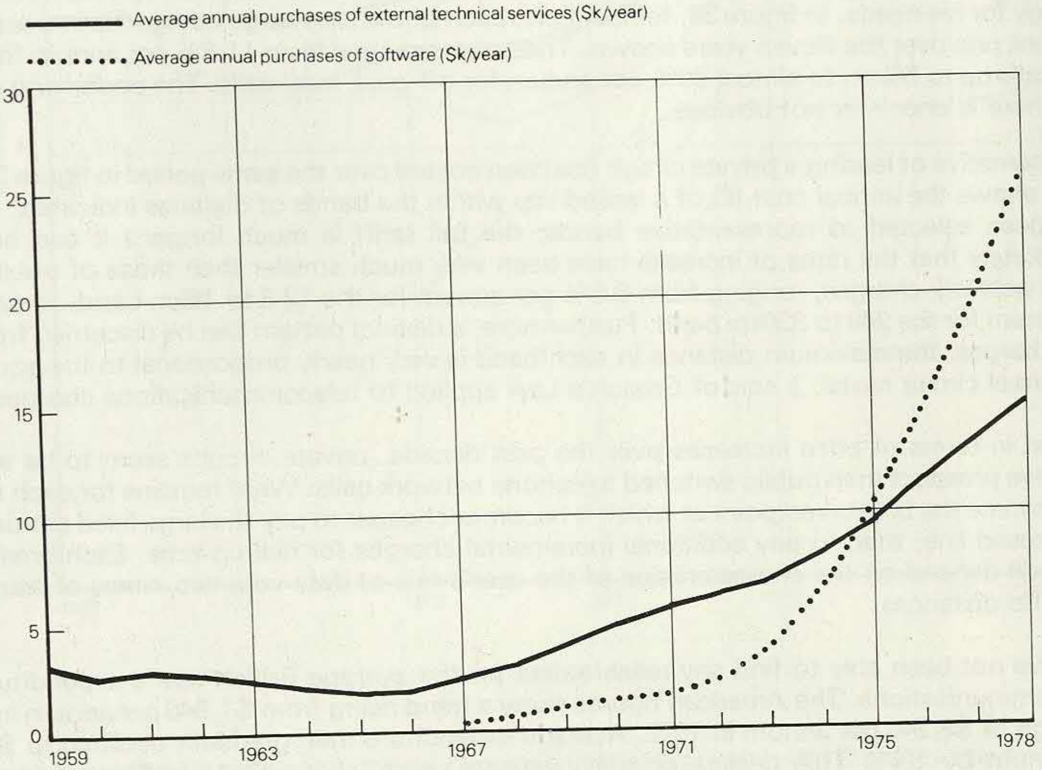

The two lines show the average annual expenditure of mainframe computer installations in the US on externally-supplied software and technical services. While the latter has increased steadily, expenditure on software has exploded since 1971 as more and more users have to pay for "unbundled" operating system software.

(Source: IDC)

DATA COMMUNICATIONS SERVICES

Since 1970, many computer users have found themselves with a completely new expense: the cost of data communications among computers and other devices located at different sites. Unless the devices are physically close enough to one another to permit them to be cabled together, this communication requires the user to engage the services of an external organisation. In virtually all European countries, this organisation is normally the national PTT.

Most users of data communications services have to choose between making dial-up connections over the public switched telephone network, and leasing one or more private lines for their exclusive use. Their decisions will depend to a large extent on the relative costs of these options.

The cost of using the public switched telephone network in the UK since 1969 is shown in figure 38. Although the cost is shown on a log scale, the increases still appear to be large, and they *are* large: the average cost of a local call lasting one minute after 1 pm grew by a factor of 17.

PTT pricing is arrived at through a number of factors (economic, political and social) and so there is little consistency in either the timing or the size of the price increases.

This erratic policy makes it difficult for the would-be user to work out an optimum long-term strategy for his needs. In figure 38, for example, each of the six charge categories increased at a different rate over the eleven years shown. These rates ranged from 11.6% per annum for peak-time calls up to 56km, to almost 30% per annum for off-peak local calls. The underlying pattern — if there is one — is not obvious.

The alternative of leasing a private circuit has been costed over the same period in figure 39. The graph shows the annual cost (£) of a leased line within the bands of distance indicated. (These have been selected as representative bands; the full tariff is much longer.) It can be seen immediately that the rates of increase have been very much smaller than those of public telephone network charges, ranging from 6.8% per annum for the 12.8 to 16km band, to only 1% per annum for the 240 to 320km band. Furthermore, a distinct pattern can be discerned from the 1980 charges: the maximum distance in each band is very nearly proportional to the square of the annual circuit rental, a sort of Grosch's Law applied to telecommunications charges.

At least in terms of price increases over the past decade, private circuits seem to be a more attractive prospect than public switched telephone network calls. What remains for each user is to determine the breakeven point at which it becomes cheaper to pay the large fixed annual cost for a leased line, than to pay additional incremental charges for dial-up time. Each breakeven point will depend on the characteristics of the user's mix of data volumes, times of transmission, and distances.

We have not been able to find any reliable data for the average British user's expenditure on data communications. The American figures show a trend rising from $1,940 per annum in 1966 to a high of $8,240 per annum in 1972. Average expenditure then gradually declined to $6,500 per annum by 1978. This decline probably reflects users' changing to in-house interactive computing in place of bureau timesharing. It may also reflect the decline in data transmission costs in the US, because some tariffs have actually been reduced over this period. Some newer, competitive services have been lower priced than those of the Bell System.

Finally, we show in figure 40 the total UK user costs of transmitting data by the cheapest means. The horizontal scale represents bits transmitted per month, while the vertical scale shows the total cost (i.e. line costs plus modem costs) in £. The calculations assume that data is transmitted on weekdays only, between 0900 and 1700, and that half the data is sent between 0900 and 1300. Each of the discontinuities in the curves represents a change to the next-fastest modem, up to a maximum of 4800 bps.

Figure 38 Cost of dialled calls on the telephone network in the UK

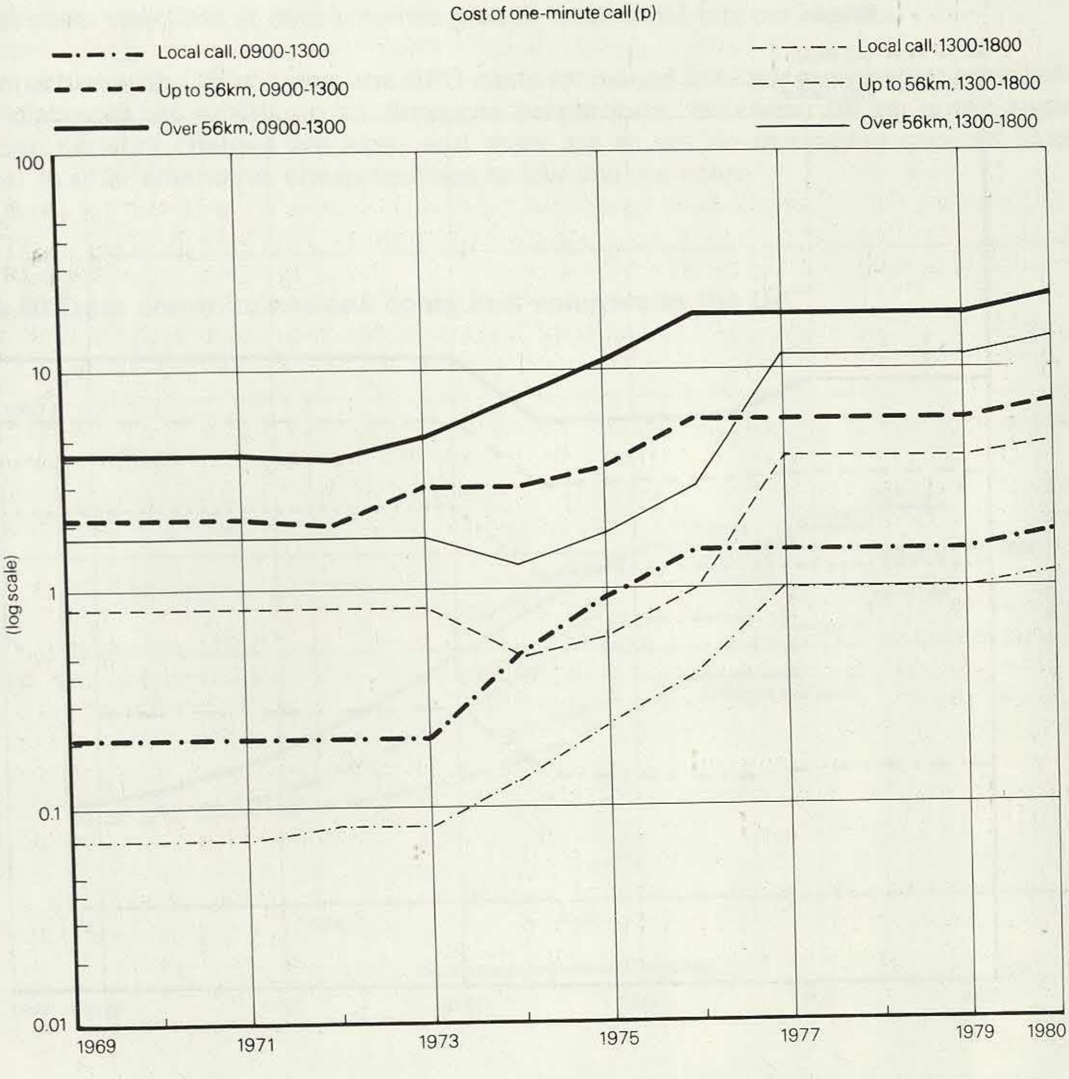

The lines show the average cost (in new pence per minute) of using the British public telephone network in the period 1969-80. The bolder lines show the costs of each of the three distance bands during the peak charging hours of 0900-1300; the thinner lines give these costs for the lower-rate hours of 1300-1800. Although the costs are plotted on a log scale, it is still clear that the increases in BPO charges have been large, uneven and unpredictable. The largest increase has been that for off-peak local calls, which have risen at a compound rate of nearly 30% p.a. over the entire period.

(Sources: *Computer Users Yearbooks,* 1969 through 1976, and BPO Dialling Codes booklets, 1977 through 1980)

Figure 39 Annual rental charges for private telephone circuits in the UK

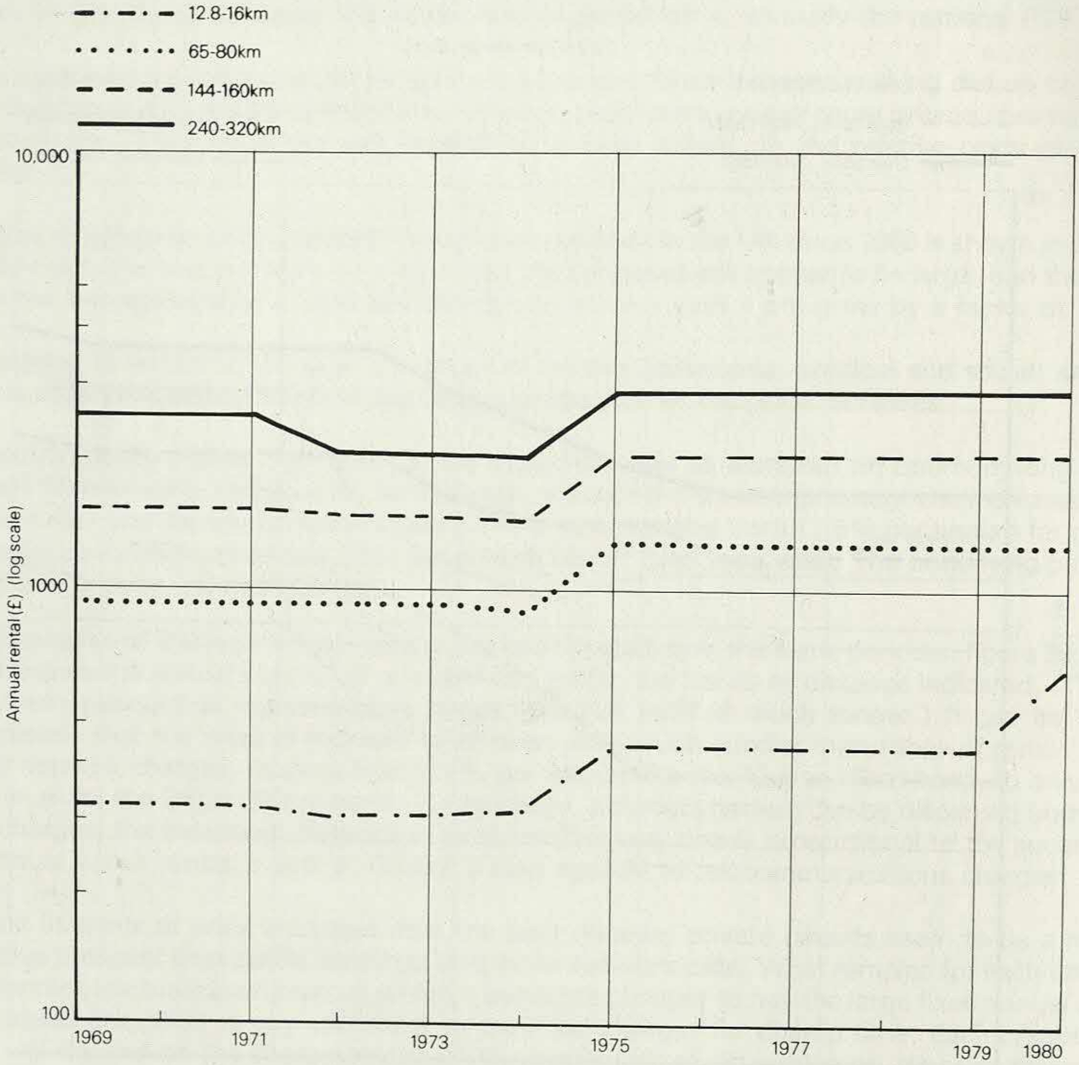

The four lines show, on a log scale, the changes in BPO charges for private telephone circuit rental in the UK, 1969-1980. Four representative distance bands are shown. The price increases have been extremely moderate compared with those of the public telephone network, as shown in figure 38. In particular, long-distance circuits (over 240km) have hardly increased in price at all over this period.

(Sources: *Computer Users Yearbooks,* 1969 through 1979, and BPO *Tariff T* (1980)).

The conclusions we draw from figure 40 are:

1. Public switched telephone network charges rise very rapidly with distance.

2. Private leased circuits are cheaper when more than 200M bits per month are to be transmitted.

3. Total costs vary little at data volumes greater than 800M bits per month.

By comparison with US charges, the BPO costs for leased lines are reasonable, provided that British distances are scaled up to American proportions. However, British public switched telephone network charges are high, and there are as yet no competing services (such as Tymnet) to offer alternative cheap facilities to low-volume users.

Figure 40 Data communications costs and volumes in the UK

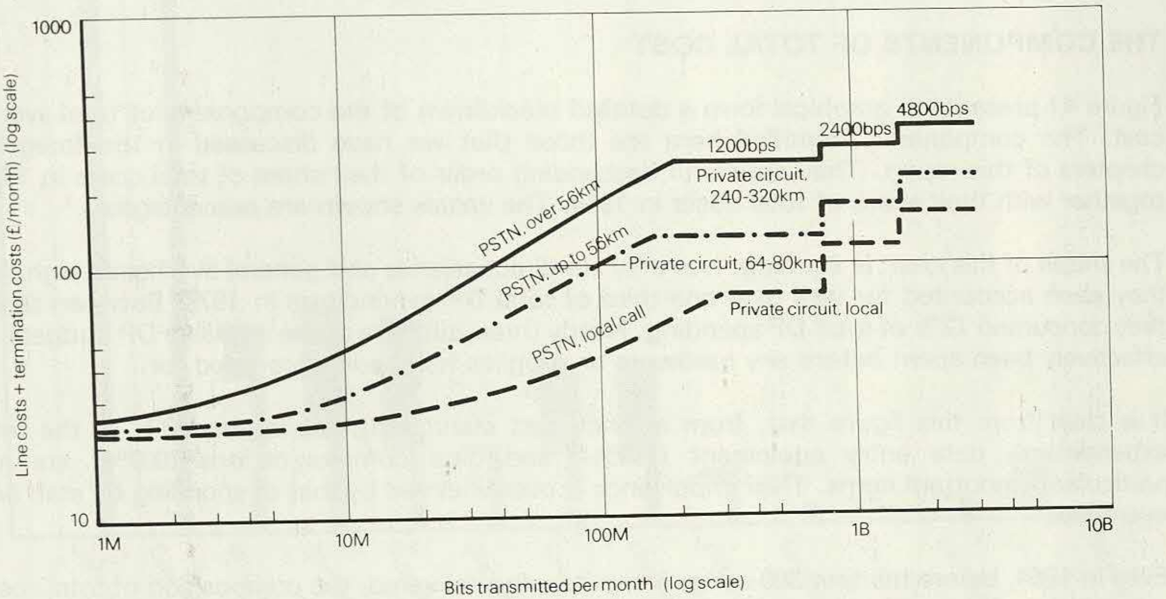

These three curves show the total monthly external costs (line plus termination costs) in the UK, 1980, for transmitting the quantity of data shown on the horizontal scale. Both scales are log. Three representative distance bands are shown: local, up to 56km, and over 56km. The curves assume the cheapest transmission methods are used (i.e. dialled calls, then private circuits with slow modems, then private circuits with fast modems). The changes to faster modems result in the step-functions in the right-hand part of the figure. It is clear that total costs increase slowly with greatly increased data volumes above about 200M bits/month.

(Source: BPO tariffs)

CHAPTER 7

TOTAL SYSTEM COSTS

Having defined and described the main trends that have influenced DP costs, we now address this chapter to three questions:

1. What relation do the components of DP costs have to each other and to total costs?

2. How are these components and total costs growing, both historically and currently?

3. What conclusions, if any, can be drawn from the evidence presented in this report?

THE COMPONENTS OF TOTAL COST

Figure 41 presents in graphical form a detailed breakdown of the components of total system cost. The components identified here are those that we have discussed in the foregoing chapters of this report. They appear in descending order of their share of total costs in 1978, together with their share of total costs in 1964. The values shown are percentages.

The shape of this chart is startling. Not only are direct salaries and general overheads high, but they *each* accounted for well over one-third of total DP expenditure in 1978. Between them, they consumed 72% of total DP spending. Nearly three-quarters of the available DP budget had effectively been spent before any hardware or supplies had been accounted for.

It is clear from this figure that, from a strict cost standpoint, processors (6% of the total expenditure), data entry equipment (1.8%), and data communications (0.9%) are not particularly important items. Their importance is overwhelmed by that of spending on staff and overheads.

Even in 1964, before the first 360-series computer was delivered, the composition of total costs was not very different from this. Salaries and overheads then totalled only 51%, but most of the difference was accounted for by dearer memory and peripherals in 1964. The balance was broadly the same then as in 1978, merely less extreme in the proportions.

From this data it can be inferred that, in 1978, salaries and wages, general overheads, and all other expenses each accounted for roughly one-third of the total spending in mainframe computer installations in the US. A little more than one-third of the total was spent on salaries and wages and on general overheads respectively and a little less than one-third on all other expenses. From this it seems that commercial computing is surprisingly labour-intensive and heavy in overheads.

GROWTH OF DP COSTS

This analysis makes it clear that expenditure on hardware has fallen from 37.8% of total DP costs in 1964 to 20.4% in 1978. However, this decline in percentage shares does not necessarily imply a decrease in actual money expenditure.

Figure 41 Components of total system costs in 1964 and 1978

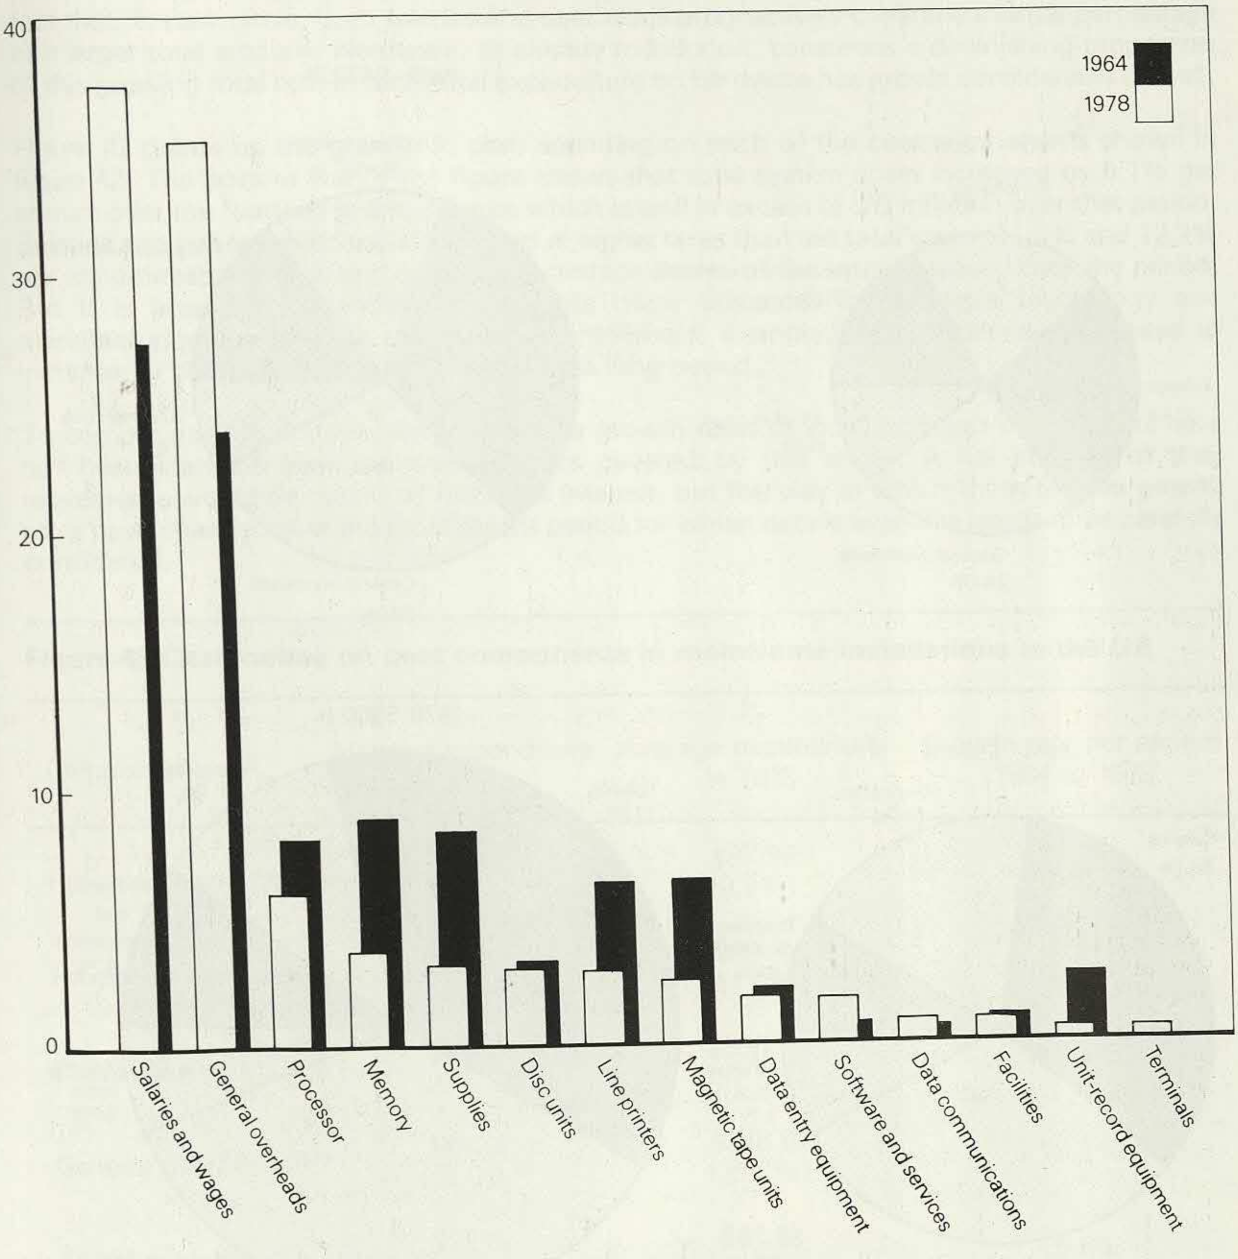

Each pair of bars shows the percentage of total system costs comprised by the cost category indicated, for the year 1978 (foreground bar) and also for 1964 (background bar). The most evident changes are the shift from hardware costs into salaries and overheads – although other figures show that most hardware costs have also increased in dollar terms.

(Source: IDC)

Figure 42 provides a quantitative picture of both the relative and the absolute growth rates of DP components, and of total DP costs. The four circles represent total DP expenditure for the years 1964, 1969, 1974, and 1978.

Figure 42 Total system costs by group, 1964 to 1978

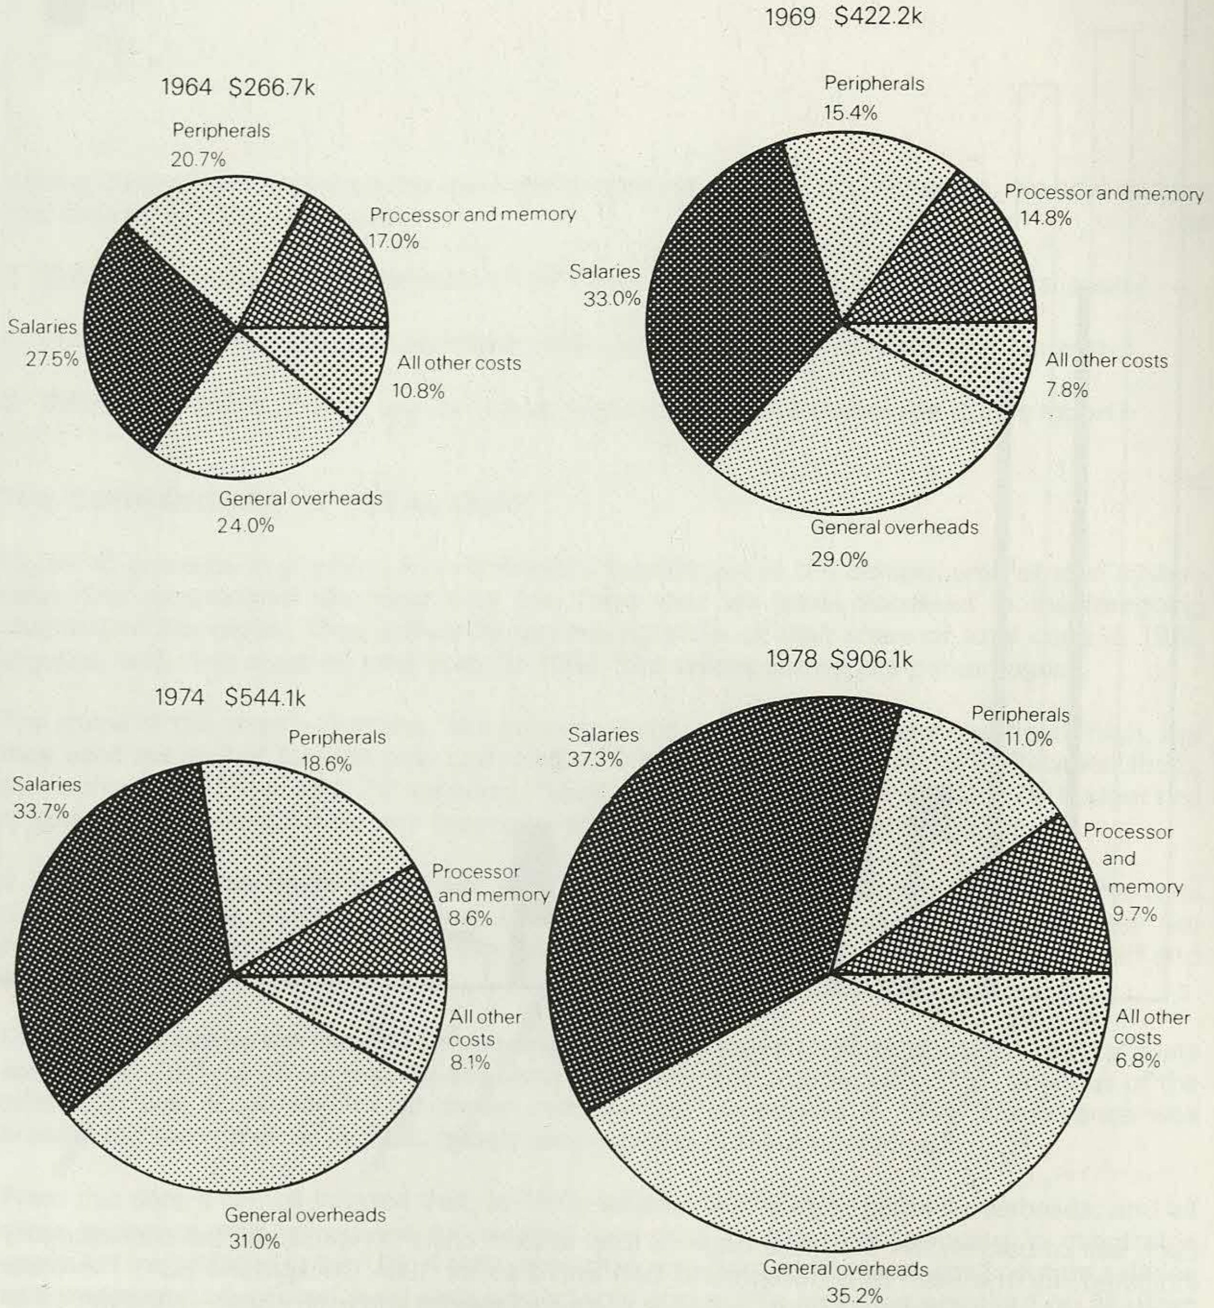

The four circles show the relative sizes of total system costs for mainframe installations in the US in the selected years 1964, 1969, 1974 and 1978. The circles are drawn to scale, and their areas show the respective total system costs. Each circle is divided into five sectors, representing the composite cost groups shown. Each cost group's percentage of total costs in that year is shown alongside.

(Sources: IDC and Phister)

The most obvious conclusions are that total DP costs have grown very rapidly over the fourteen years depicted, and that much of this growth has occurred during the most recent interval, 1974 to 1978. The growing predominance of salaries and overheads is also very apparent from the fact that, in each circle, both salaries and overheads progressively consume a larger percentage of a larger total amount. Hardware, as already mentioned, consumes a diminishing proportion of this growing total but, in fact, total expenditure on hardware has grown considerably as well.

Figure 43 points up the growth in cash spending on each of the cost components shown in figure 42. The bottom line of the figure shows that total system costs increased by 9.1% per annum over the fourteen years, a figure which is well in excess of US inflation over that period. Salaries and general overheads increased at higher rates than did total costs (11.5% and 12.2% per annum respectively), and so their percentage shares of the total increased over the period. But it is interesting to note that, despite many advances in hardware technology and manufacturing, users' cash spending on processors, memory, and peripherals continued to increase by about 4.5% per annum over this long period.

To complicate this picture even further, the growth rates of individual cost components have not been identical over the twenty years covered by this study. A full analysis of their movements would be mainly of historical interest, but the way in which these relative growth rates have changed over the most recent period for which data is available needs to be carefully considered.

Figure 43 Cash outlay on cost components in mainframe installations in the US

Component	Average expenditure in 1964	Average expenditure in 1978	Growth rate per annum (1964 to 1978)
Processors and memory	$45.3k (17.0%)	$87.9k (9.7%)	4.8%
Peripheral equipment	$55.2k (20.7%)	$99.7k (11.0%)	4.3%
Direct salaries	$73.4k (27.5%)	$338.0k (37.3%)	11.5%
General overheads	$64.0k (24.0%)	$318.9k (35.2%)	12.2%
All other costs	$28.8k (10.8%)	$61.6k (6.8%)	5.6%
Total costs	$266.7k (100%)	$906.1k (100%)	9.1%

(Source: Phister)

Figure 44 presents the growth rates from 1974 to 1978 of each of the cost components covered by this report. In theory, these growth rates should give an approximate idea of how rapidly each category of cost is changing today.

Figure 44 Recent growth of system cost components

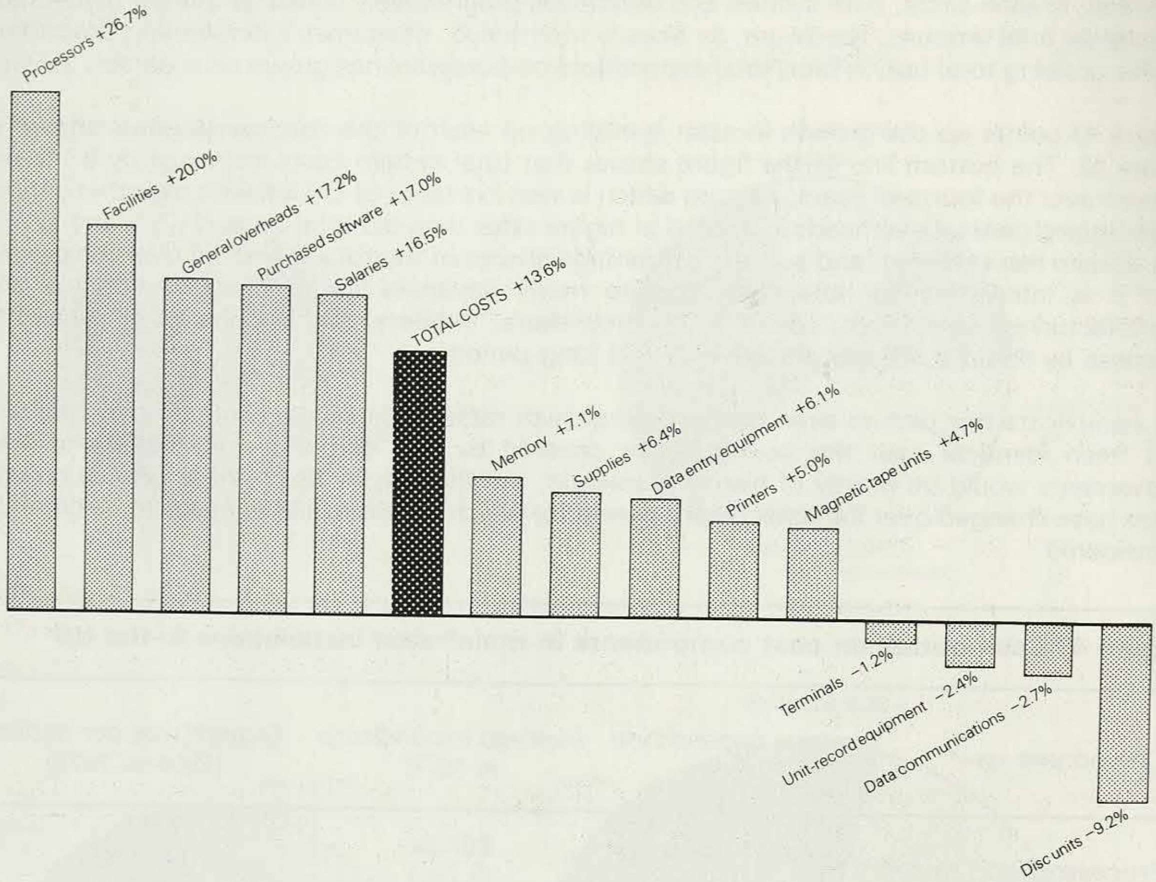

The bars in this figure show the average annual compound change in the value of each of the cost components identified in this report. The period covered is the recent five-year interval 1974-78. One surprising result is that the largest percentage increase has been registered by processors, at 26.7% p.a. growth in expenditure by mainframe users in the US. The average increase in total DP costs over this period was 13.6% p.a. The cost trends of some components shown here are quite different from their long-term trends and possibly are of more direct relevance to users.

(Source: IDC)

What the figure shows is that the most rapidly growing category of expenditure is processors, at 26.7% per annum, followed by facilities at 20%, and so on down the list to discs, which actually showed a 9.2% per annum decline in user expenditure over the period. The contrast between processors and discs is an important one because, as discussed in Chapters 2 and 3, progress has been steady and rapid in both technologies. The two patterns of expenditure might reasonably be expected to be roughly similar, but in fact they fall at opposite extremes of the spectrum.

Total mainframe users' costs between 1974 and 1978 grew at 13.6% per annum, which was well in excess of the increase in the US price index over the same period.

THE CONCLUSIONS THAT WE DRAW

Most of the information presented in this report is quantitative and historical. It has provided a fairly detailed picture of how computer users, on average, have spent their money on various components of cost. It has also identified the most important trends that have affected these costs.

We now draw some important non-quantitative conclusions based on both an analysis of this data and our own observations.

1. Labour-intensive computing
During the period 1964 to 1978, the proportion of the average data processing budget consumed by salaries and direct overheads increased from 51% to 72%. This did not happen merely because (as some might think) computer equipment has reduced in cost. Although equipment *has* reduced in cost, users' total expenditure on mainframe equipment has nevertheless continued to rise, and it is currently rising sharply.

   The fact that data processing is a labour-intensive activity leads rapidly to two conclusions:

   - A very high priority must be placed on the search for ways of making major staff productivity gains in computer installations.
   - Users must ensure that their working methods reflect the high cost of staff. This requires that people who have been involved with computing since the time when the efficient use of computers (at the expense of people) was the accepted norm, now need to re-appraise their attitudes and assumptions.

2. Computer power and hardware cost
There is vastly more computer power in installations today than there was ten or fifteen years ago. In spite of its declining incremental cost, users spend more today for computing than ever. We consider that there are two explanations for this:

   - Because the use of computing grew rapidly, with little time for consolidation, many application portfolios are a hodge-podge of uncoordinated programs, redundant files, and special operating instructions. This situation gives rise to inefficient computing, not merely in terms of wasted machine cycles, but also in terms of the misuse of human resources. Operating system software has also become vastly more complex, and it requires much larger hardware configurations and more staff to support it. This loss of efficiency (and possibly of effectiveness) may well be a price that had to be paid for quickly harnessing a fast-emerging technology with few guidelines to go by. But it must change if total system costs are to be controlled.
   - Most organisations, when embarking on computerisation, first selected applications that appeared to be easy. Thus the "obvious" applications, such as payroll, accounting, and sales order processing, were computerised long ago. What was left were systems that were more complex and more difficult, were sometimes only partially understood by systems staff, and perhaps were only marginally justifiable from a cost/benefit standpoint. These are just the kinds of systems that most often result in late projects and costly over-runs. Additional expenditure on computing, much of it unplanned, frequently resulted.

3. Computer hardware costs
Every kind of hardware component has declined in unit cost per constant function. The costs of memories, discs and processors have declined most rapidly of all. Yet most users' applications are still biased towards an obsolete view of hardware costs. This view assumes that the processor is the most costly resource, and it also assumes that memory needs to be

shared out "efficiently". Consequently, users have found it necessary to add considerable additional hardware in order to accommodate only a modest number of additional applications.

This fact partially explains the runaway success of the minicomputer. When stripped of all the associated complexities and operating system constraints, an applications program (particularly an interactive one) can frequently be made to run comfortably on a small stand-alone computer. The costs of the extra hardware may be only marginally more than the true cost of developing the same application on a mainframe computer, with its traditionally long implementation time and complex interactions with other jobs.

On mainframe computers, hardware costs seem to increase faster than the rate at which additional applications can be added to the existing workload.

4. People costs
Contrary to a widespread belief, it appears that there has been no recent explosion in the salaries of DP staff. Salaries are indeed a growing component of total costs, but only because salaries have increased steadily with general wage inflation, whilst hardware costs have changed erratically over the years. Currently, hardware costs are increasing at a faster rate than either salaries or inflation.

There seems to be a strong tendency for DP organisations to grow in numbers of DP staff, and this is where most of the real growth in staff costs originates. Some of this growth has certainly been due to the ever-wider range of applications that DP departments have undertaken and also to the growing complexity of software. The rest is not so easily accountable for, and it may be attributable to inadequate training and standards.

5. Overhead costs
Expenditure on staff-related general overheads can only be described as lavish. It consumed 35% of all DP expenditure in the US in 1978, and was still growing rapidly.

The most disturbing fact about overheads is that they are growing as a percentage of direct salaries and wages. This implies that more and more support staff are required for each member of the DP staff. Although there is no obvious reason why this should be so, mainframe users clearly are finding that it is.

Given the very low prices of some kinds of hardware (particularly terminals and microcomputers) we believe that more and more users will find ways of trading off overhead costs for expenditure on new and novel items of hardware for their DP staffs. They may then leave it to the ingenuity of their staffs to find useful internal applications for this hardware, offsetting the absence of the overhead item(s) they have sacrificed. Internal electronic mail systems are an example of this tendency.

6. Manufacturers and the market
The sectors of the hardware market that have shown the fastest reductions in price are those in which users could choose freely from a variety of offerings. The manufacturers of plug-compatible memories, discs, and processors have done the market an enormous service by forcing the major manufacturers to compete on price.

However, there is a strong movement today in every sector of the market to sell "complete systems" to end users. These systems typically return a higher unit profit to their manufacturers than would selling the individual components in a competitive market. Invariably, some of the components in such a system are overpriced compared with what the user could buy on the open market if he wished to.

However, either for inner feelings of security or for some other reason, many users continue to buy such packaged systems unquestioningly from their traditional suppliers. This means that many users are not reaping the full benefit of the drop in hardware costs over the last ten years. It also means that their system concepts tend to be confined to the offerings of just one or two suppliers. Most of these offerings seem to be designed primarily to sell additional hardware.

CHAPTER 8

CONCLUSION

In describing the main trends in the costs of, and expenditure on, data processing, we have also resolved the apparent paradox mentioned in Chapter 1. Users *are* spending more on computing today than they were in 1960, despite major falls in the unit costs of virtually all hardware. The reason is that today's computer installations are massive compared with those of twenty years ago. (The reductions in physical size of today's processors tend to mask this fact.)

Mainframe computers are, on average, several hundred times more powerful than their counterparts of 1960. Main memories are hundreds of times larger. Storage devices have hundreds of times the capacity. So, while users have spent vastly more on staff and overheads, their increased expenditure on hardware seems to have been well repaid in terms of the additional capacity it has bought.

However, the other side of this equation cannot be taken for granted. If today's computers are hundreds of times as powerful, are they also hundreds of times as useful to their owners? If the results produced by computers are not much more useful today, then all the "progress" of the last twenty years is at best irrelevant, and at worst is a distracting side-show, to the commerical users of computers.

There is a "law" (attributed to M. Amaya) which states: "Through several generations of computers, the average execution time of a given job is independent of the speed of the computer." Like most such "laws", this one is intended partly to amuse with its cynicism, and partly to express an empirical result experienced by users. In fact, it seems to be supported by a certain amount of statistical evidence, although we have found no serious attempts to explain it quantitatively.

A future Foundation report will examine the issue of value-for-money in expenditure on computing, and will seek to discover whether a commensurate benefit to the organisation has accompanied the increase in spending.

BIBLIOGRAPHY

1. Anon. *IBM: The Plug Compatible Marketplace.* Online Publications Ltd., 1978.

2. Auerbach Corporation. Various publications and reports.

3. Brewer, J. A. "Improving Systems' Productivity". The Butler Cox Foundation Report Series No. 11, February 1979.

4. British Post Office. *Tariff T,* 1980.

5. Champine, George A. *Computer Technology Impact on Management.* New York: North-Holland Publishing Company, 1978.

6. Computer Users Yearbooks, 1969-1979.

7. Dolotta, T. A., et al. *Data Processing in 1980-1985: A Study of Potential Limitations to Progress.* New York: John Wiley & Sons, 1976.

8. International Data Corporation. Various publications and reports.

9. Kelly, P. T. F. "New non-voice customer services and terminals using digital techniques". British Post Office, 1979.

10. Nyborg, Philip S., Pender McCarter and William Erickson (eds). *Information Processing in the United States: A Quantitative Summary.* Montvale, New Jersey: AFIPS Press, 1977.

11. Phister, Montgomery. *Data Processing Technology and Economics.* Santa Monica: Santa Monica Publishing Company, 1976.

12. Phister, Montgomery. *Data Processing Technology and Economics, 1975-1978 Supplement.* Santa Monica: Santa Monica Publishing Company, 1979.

13. Withington, Frederic G., and Oscar H. Rothenbuecher. *The World Computer Industry, 1976-81.* Cambridge, Mass.: Arthur D. Little, Inc., 1976.

Photoset and printed by Flexiprint Ltd, Worthing, West Sussex

The Butler Cox Foundation

Butler Cox & Partners Limited
Morley House, 26-30 Holborn Viaduct, London EC1A 2BP
Tel 01-583 9381, Telex 8813717 GARFLD

Butler Cox & Partners Limited
216 Cooper Center, Pennsauken, New Jersey 08109, USA
Tel (609) 665 3210

Italy
Sisdoconsult
20123 Milano – Via Caradosso 7 – Italy
Tel 86.53.55/87.62.27

Benelux, Germany, Switzerland
Akzo Systems bv
Velperweg 76, Arnhem, The Netherlands
Tel 85-662629

France
Akzo Systems France
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
Tel (1) 820.61.64

The Nordic Region
Statskonsult
PO Box 4040, S-171 04 Solna, Sweden
Tel 08-730 03 00, Telex 127 54 SINTAB-S

Printed in England